SEVENTH FRAMEWORK PROGRAMME	**SEVENTH FRAMEWORK PROGRAMME** **Trustworthy ICT**

Project Title:

Enhanced Network Security for Seamless Service Provisioning in the Smart Mobile Ecosystem

Grant Agreement No: 317888, Specific Targeted Research Project (STREP)

DELIVERABLE

D4.3 Anomaly detection within femtocell architectures

Deliverable No.	**D4.3**		
Workpackage No.	**WP4**	Workpackage Title	**Anomaly Detection Using Control Plane Data**
Task No.	**T4.3**	Task Title	**Anomaly detection within femtocell architectures**
Lead Beneficiary	**TUB**		
Dissemination Level	**PU**		
Nature of Deliverable	**R**		
Delivery Date	**20 November 2014**		
Status	**F**		
File name	**NEMESYS_Deliverable_D4.3.pdf**		
Project Start Date	**01 November 2012**		
Project Duration	**36 Months**		

Authors List

Author's Name	Partner	Email Address
Lead Author / Editor		
Steffen Liebergeld	TUB	steffen@sec.t-labs.tu-berlin.de
Co-Authors		
Matthias Lange	TUB	mlange@sec.t-labs.tu-berlin.de
Ravishankar Borgaonkar	TUB	ravii@sec.t-labs.tu-berlin.de
Anastasios Drosou	CERTH	drosou@iti.gr
Vasilios Mavroudis	CERTH	mavroudis@iti.gr
Gokce Gorbil	ICL	g.gorbil@imperial.ac.uk
Omer H. Abdelrahman	ICL	o.abd06@imperial.ac.uk
George Lyberopoulos	COSMOTE	glimperop@cosmote.gr
Eleni Theodoropoulou	COSMOTE	etheodorop@cosmote.gr
Konstantinos Filis	COSMOTE	cfilis@cosmote.gr
Ioanna Mesogiti	COSMOTE	imesogiti@cosmote.gr
Eftychia Nikolitsa	COSMOTE	enikolitsa@cosmote.gr
Madalina Baltatu	TI-IT	madalina.baltatu@it.telecomitalia.it

Reviewers List

Reviewer's Name	Partner	Email Address
Roberta D'Amico	TI-IT	roberta.damico@it.telecomitalia.it

Contents

List of Figures

Abbreviations

3GPP	Third Generation Partnership Project
Auc	Authentication Center
BRPCA	Bayesian Robust Principal Component Analysis
CDMA	Code Division Multiple Access
CN	Core Network
CSG	Closed Subscriber Group
CUSUM	CUmulative SUM
DNS	Domain Name Server
DoS	Denial of Service
DSL	Digital Subscriber Line
FAP	Femtocell Access Point
FGW	Femtocell GateWay
FMC	Fixed Mobile Convergence
GAN	Generic Access Networks
GANC	Generic Access Networks Controller
GSM	Global System for Mobile Communications
HLR	Home Location Register
HMM	Hidden Markov Model
HSS	Home Subscriber Server
HNB-GW	Home Node-B Gateway
IMSI	International Mobile Subscriber Identity
ISP	Internet Service Provider
L3	3GPP Layer 3
LOF	Local Outlier Factor
LTE	Long Term Evolution
MNO	Mobile Network Operator
MSISDN	Mobile Station International ISDN Number
NDS	Network Domain Security
NMS	Network Management System
OAM	Operation, Administration, and Management
P2P	Peer to Peer
RNC	Radio Network Controller
SEGW	SEcure GateWay
SGSN	Serving GPRS Support Node
SMSC	Short Message Service Center
SRTP	Secure Real-Time Transport Protocol

TMSI	Temporary Mobile Subscriber Identity
UE	User End-point
VLR	Visitors Location Register
VoIP	Voice over Internet Protocol

Abstract

Femtocells are small, low-power cellular base stations which are typically designed for use in a home or small business environment. Although the deployment of femtocells has advantages for both the mobile network operator and the end-user, such as improved network coverage and capacity indoors, femtocells create new security vulnerabilities and potential attack vectors that need to be identified and addressed. In this document, we first investigate the vulnerabilities of, and attacks specific to femtocells, and then propose a femtocell-based honeypot architecture called the *cellpot* in order to monitor the femtocells deployed within a mobile network, and later to mitigate against attacks originating from femtocells. Our proposed cellpot concept protects the core network from compromised femtocells by separating the firmware and the infrastructure environments within the device by virtualization technology. Every cellpot-enabled femtocell device joins a peer-to-peer network which enables the distributed sharing of information among the cellpots, allowing them to inform each other of any detected suspicious activity and paving the way for smart mitigation against attacks originating from femtocells.

We also investigate and evaluate several anomaly detection methods which can be used to identify anomalies and malicious activities within femtocells, using realistic data traces from simulated mobile networks with compromised femtocells. The methods that we investigate are based on the cumulative sum (CUSUM), local outlier factor (LOF), hidden Markov models (HMM), and Bayesian robust principal component analysis (BRPCA). Our evaluation shows that while all of the considered methods are able to detect the investigated anomalies within femtocells, CUSUM is the least stable and accurate, while HMM and BRPCA are the most stable and have consistently good performance. Since BRPCA does not require supervised training and is computationally more than two times faster than HMM, we believe that BRPCA is the best anomaly detection method for the identification of anomalies in femtocells in large mobile networks.

1 Introduction

Femtocells are small, low-power cellular base stations which are typically designed for use in a home or small business environment. Mobile networks operators (MNOs) deploy femtocells in order to improve network coverage and capacity, especially indoors. However, the use of femtocells create new security vulnerabilities and potential attack vectors. For example, due to easy physical access to the device, attackers are better able to compromise femtocells and weaponize them to attack cellular networks. Therefore, MNOs need to understand the additional vulnerabilities due to femtocell architectures, detect attacks originating from femtocells, and mitigate against such attacks in order to protect their infrastructure and services.

In this document, we first review the vulnerabilities of femtocells, and propose a femtocell-based honeypot framework, called the *cellpot*, in order to address attacks originating from femtocells. Honeypots have proved to be valuable in the detection and analysis of network attacks in the Internet, and we have proposed a smartphone-based honeypot in order to extend the concept to mobile networks in the NEMESYS project. In order to improve the coverage of our smartphone-based mobile honeypot, and to monitor and protect femtocells deployed in a mobile network, we propose the deployment of a femtocell-based honeypot framework called the cellpot, which allows the threat identification, anomaly detection and defense mechanisms to operate within the mobile network under the control of the operator. Cellpots are a cost-effective and scalable method to detect anomalies and attacks originating from femtocells, and are ideally situated to enable the operator to mitigate against such attacks once they have been identified, for example by reducing the signaling overhead and thereby protecting the core network, which in effect allows the operator to continue to provide uninterrupted and high quality mobile services to its customers. We discuss how cellpots can be utilized to protect the mobile network against common threats and attacks, such SMS spam, mobile theft, and mobile malware.

We also investigate several anomaly detection methods in order to identify anomalies in femtocells; the methods we consider are based on the cumulative sum (CUSUM), local outlier factor (LOF), hidden Markov models (HMM), and Bayesian robust principal component analysis (BRPCA). We evaluate these methods using realistic data traces from simulated mobile networks with compromised femtocells. In our evaluation, we consider

currently the most common mobile malware, premium SMS messages, where an attacker has compromised one or more femtocells that generate premium SMS messages on behalf of the mobile devices connected to the malicious femtocell. Our evaluation shows that while all of the considered methods are able to detect the investigated anomalies within femtocells, CUSUM is the least stable and accurate although it is also the most computationally efficient. We observe that HMM and BRPCA are the best methods in terms of stability and detection rate. However, we believe that BRPCA outperforms HMM since it does not require supervised training and it is computationally more than two times faster than HMM. We therefore recommend BRPCA as the anomaly detection method of choice for the identification of anomalies in femtocells in large mobile networks.

The rest of this document is structured as follows. Chapter 2 presents background material in the area of security of femtocells. In Chp. 3, we present a review of several femtocell-specific attacks which affect the mobile customers and the core network, and also describe the threat model that we assume. Our cellpot concept is discussed in Chp. 4. Chapter 5 presents the anomaly detection algorithms we have considered and their evaluation in detecting attacks originating from femtocells. We conclude by providing a summary of our main findings in Chp. 6.

2 Background

2.1 Femtocell architectures

A femtocell is a small cellular base station, typically designed for use in a home or small business environment. The femtocell base station connects to the service provider's network via broadband, e.g. DSL or cable, which is shared publicly. Current designs support 2–4 or 8–16 active mobile phones if used in a residential or enterprise setting, respectively. Femtocells allow service providers to extend service coverage indoors, especially in cases where access would otherwise be limited or unavailable. Although most of existing femtocell deployments are in 3G systems, the concept is applicable to all standards, including GSM, CDMA2000, TD-SCDMA, WiMAX, and LTE. Femtocells practically offer local mobile cell sites to customers, enhancing the capabilities of the mobile network exactly at the points where the service demand is highest, such as inside homes or business premises, and support the same radio standards as macrocell base stations, thus enabling seamless handovers between femtocells and macrocells.

A typical femtocell solution is presented in Fig. 2.1, which is composed of:

- the femtocell access point (FAP), which provides the radio interface,
- the femtocell gateway (FGW), which manages the access point and authenticates communications, acting as the security gateway,
- an authentication, authorization, and accounting (AAA) server,
- a DNS/DHCP/NTP server,
- a set switches and routers that interconnect the femtocell nodes and the femtocell platform to the rest of the operator's core network, and
- a network management system (NMS).

Femtocells operate in a licensed spectrum and they are designed to route mobile phone traffic through a home or corporate IP network. A femtocell provides voice and data services to connected mobile devices. The range covered is very limited, normally within

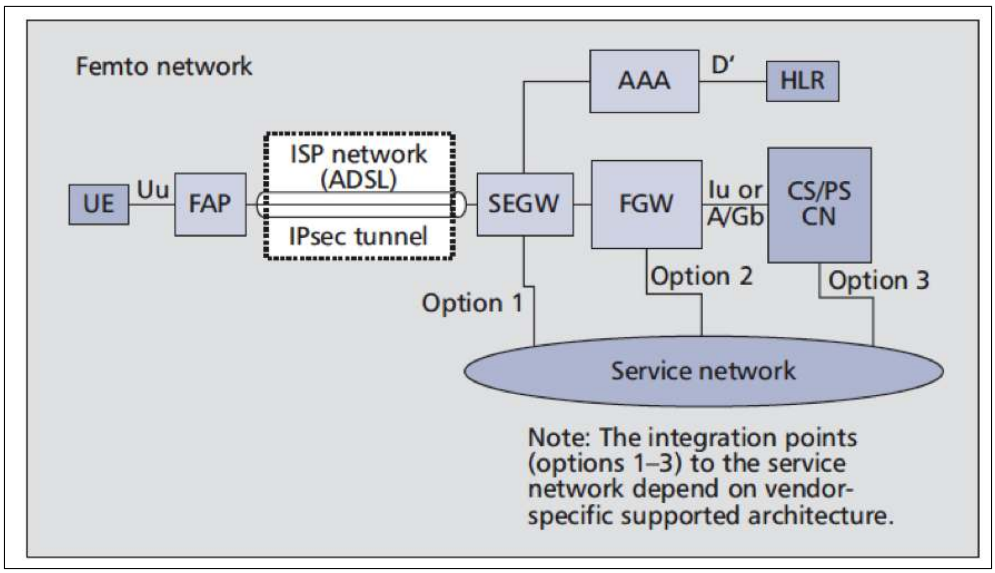

Figure 2.1: A typical femtocell solution. The femtocell access point (FAP) provides the radio interface to the user equipment (UE). The FAP is normally connected via public broadband to a security gateway (SEGW), which authenticates the communications to/from the femtocell. The femtocell gateway (FGW) manages the FAP and may also incorporate the SEGW functionality into a single device. The authentication, authorization, and accounting (AAA) server provides AAA services to the femtocell architecture with the help of the HLR. The femtocell architecture can be connected to the operator's service network via alternative points as shown in the figure.

a house or a small business. The units are designed to be plug-and-play for easy installation. The main features of femtocells include automatic detection of the ISP, automatic registration, authentication to the mobile network, self-upgrades, location check, and transmit power adjustment. Femtocells should not be confused with repeaters, also known as signal boosters, which are used to improve existing macrocell coverage but do not provide any base station functionalities.

Advantages of femtocells

For a mobile operator, the benefits that femtocells offer are improvements to both coverage and capacity, especially indoors. This can reduce both capital expenditures and operating expense. Subscriber satisfaction is greatly improved, since customers benefit from improved coverage and potentially better voice quality and increased battery life. Subscribers may also be offered more attractive tariffs, e.g. discounted calls from home.

Femtocells are sold by an operator to its residential or enterprise customers. A femtocell is typically the size of a residential gateway or smaller, and connects to the user's broadband line. Integrated femtocells (which include both a DSL router and femtocell) also exist. Once plugged in, the femtocell connects to the MNOs mobile network, and provides indoor coverage in a range of typically 30 to 50 meters (line of sight) for residential premises. The output power of a femtocell is usually 20 mW, which is five times less than a Wifi access point. From a user's perspective, there is no specific installation or technical knowledge required.

Once installed in a specific location, most femtocells have protection mechanisms so that a location change will be reported to the MNO. Whether the MNO allows femtocells to operate in a different location depends on the MNOs policy. International location change of a femtocell is not permitted because the femtocell transmits licensed frequencies which belong to different network operators in different countries.

For MNOs, femtocells may constitute a solution to:

- Increase mobile usage indoors, and thus revenues, by combining coverage/capacity enhancements with inexpensive voice services,
- Offer new, innovative data services (music/photo/video download synchronization, mobile TV, etc.), thus making the mobile phone competitive to a fixed phone, PC and/or TV, and
- Offer fixed mobile convergence (FMC) in response to WiFi, VoIP, Homezone and unlicensed mobile access (UMA) offerings.

In addition, femtocells may contribute to:

- Churn reduction (e.g. by "capturing" all the members of a family),
- OPEX savings on (macro) backhaul network (due to traffic offload),
- CAPEX savings since no new macrocell base stations or capacity expansions are needed.

However, prior to the commercial introduction of femtocells, an operator has to address a considerable list of issues, including possible interference with macrocells, impact on the core network, security concerns, interoperability, regulatory/EMC concerns, use of SLAs for QoS guarantees (especially for voice) over the broadband connection, availability of platforms/features, and packages to be offered to the customers.

The main benefits for an end-user are the following:

- Improved indoor coverage when there is no existing signal or poor coverage,
- Higher mobile data capacity,
- Special tariffs within the femtocell area,
- "One phone, one number, one bill" and converged services both indoors and outdoors.

2.2 Security architecture of femtocells

The home Node-B (HNB) is installed in the users' premises and its traffic is tunneled via public Internet (wired broadband) connection. Hence for the MNO, it is important to ensure that the HNB protects the communication over the insecure public Internet, and over the air-link between itself and the mobile device. Figure 2.2 shows the HNB security architecture, which is described in [1]. The main components of the femtocell security architecture are the HNB device, security gateway (SeGW), user equipment (UE), and the OAM (Operation, Administration, and Management) server. These components are described below.

Home Node-B The main function of the HNB is to act as a small base station. The HNB connects the UE via its radio interface to the mobile service operator's core network. It transmits the UE data by establishing a secure IPsec [16] ESP tunnel with the SeGW over an insecure backhaul link (broadband Internet connection). Moreover, it supports a set of functions provided by the radio network controller (RNC) in 3G networks, mutual authentication between the UE and the SeGW, and provides standard encryption

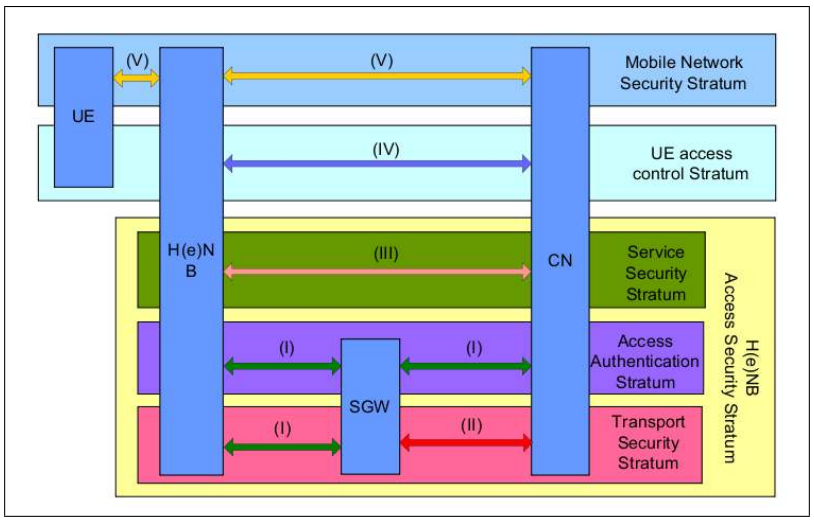

Figure 2.2: The security architecture of femtocells

mechanisms over the insecure radio link. It communicates to the OAM directly using a secure link or through the SeGW.

Security gateway The SeGW acts as a border gateway of the operator's core network. First, it mutually authenticates and registers the HNB to establish a secure IPsec tunnel, and then forwards all the signaling and the user data to the operator's core network. Mutual authentication can be performed using certificates. The interface between the SeGW and the operator's core network is considered to be secured. The SeGW accepts traffic only from the HNB and discards unwanted malicious traffic at the border.

HNB management system The HNB management system is a management server that is responsible for the configuration and the provisioning of the user data according to the operator's policy. It can be functioned to provide the required software updates on the HNB and can be located inside the operator's core network. However it is considered to be in an insecure domain if located outside of the operator's core network.

AAA server and HSS The subscription data and authentication information is stored in the HSS. The AAA server authenticates the hosting party (the HNB) by accessing the authentication information from the HSS. Both the AAA server and the HSS are deployed in the operator's core network.

HNB gateway The HNB gateway performs the access control for the non-CSG (Closed Subscriber Group) capable UE attempting to access a HNB. The SeGW can be integrated with a HNB-GW, and if not integrated then, the interface between SeGW and HNB-GW may be protected using NDS/IP (Network Domain Security/IP network layer security).

UE The UE is a standard user equipment that supports the 3G (UMTS) communication. It connects to the HNB over-the-air using a 3G AKA (Authentication and Key Agreement) procedure.

In the following sections of this document, we will be referring to the femtocell as the FAP, HNB or small cell interchangeably.

2.3 Security requirements in femtocell architectures

When an important component of wireless system is located at a customer premise, such as the Femto Access Point (FAP), the convenience of the equipment location is tempting enough to attract attackers and hackers both occasional and professional. Due to the fact that a number of femtocell systems had been developed and deployed as a result of speedy need to reach market quickly, these systems were based on some of the older security assumptions and technologies and therefore experienced the most scrutiny from the hacker community. These attacks and other common threats make it imperative that such a Femto system is designed with stringent security requirements from the very beginning.

Security for femtocell networks spans several distinct requirements. The service provider must authenticate users as they arrive on the network. The RF link between the handset and the femtocell must be secured for both user and control plane traffic. The mobile network traffic must be placed into a virtual private network as it traverses the wired ISP network to ensure that the traffic is protected while transiting this public network and only authorized users can forward traffic to the mobile operator's network. The first two elements of the security equation are specified by the existing mobile network standard (i.e. GSM, UMTS, etc.) as the handset will interact with the femtocell as if it were a macrocell. The use model for the VPN established between the femtocell and the carrier network has been defined by ETSI and is based upon the well known IPsec standard. Lastly, there is also the requirement to support voice-over-IP or SIP security which is governed by the IETF standard known as SRTP (Secure Real-time Transport Protocol). Therefore, the solution set required for a femtocell is a complex amalgam of well known security standards knitted into a comprehensive solution.

Femtocell solutions necessitate the establishment of IPsec tunnels between FAP-FGW, over which traffic/signalling/OAM traffic is encrypted, while the IKEv2 and IPsec EAP protocols offer confidentiality. All vendors support normal core network authentication procedures between the UEs and the MSC/HLR. Air-interface ciphering and integrity protection is not supported by all FAP manufacturers, but the imperative to utilise ordinary handsets and UE SIMs for accessing the FAPs forces them to use ordinary UMTS (Kasumi-based) encryption and integrity protection algorithms UEA1 and UIA1, respectively.

Access to FAPs can be restricted utilizing the closed-mode, where selected users/MSISDNs per FAP can be serviced and/or group-mode where only selected users/MSISDNs may access a certain FAP-group. Closed/Group modes are not supported by all vendors. Mutual authentication/certification between FAPs-FGW can be based on: dedicated FAP SIM/USIM (EAP-SIM/EAP-AKA) , hard-coded authentication chips built-in the FAP, digital certificates stored in FAPs SIM, software coded authentication certificates pre-stored in FAPs Operating System or MAC addresses. Depending on FAPs capability (embedded SIM, built-in chip, etc.) the operator may be forced to support more than one authentication/certification options.

There is however one very important element of femtocell security which makes the implementation significantly more complex. This relates to the additional latency introduced in mobile communications by the security architecture, which must be carefully managed especially for real-time and interactive applications such as VoIP and gaming. Compounding this challenge is the unknown nature of the latency across the ISP network, which has resulted in service providers requiring latency in the femtocell to be minimized. As a result of this stringent requirement, SoC designers are adopting sophisticated traffic management features in the femtocell SoC and software to meet the latency requirements.

2.4 Honeypots

Honeypots are a well established tool for collecting intelligence about threats in IP networks. For small cells, a new form of honeypot is needed. To that end we introduce Cellpot, a novel honeypot concept to detect, collect intelligence and mitigate threats against the cellular network directly on the base station. Our concept largely avoids the costs involved with certification and validation with respect to the radio network. Recent work of Golde et al. [11] shows that the current femtocell hardware can be turned into a monitoring node within the cellular network. In this document, we present a practical software design that ensures security of the core network and the honeypot.

2.5 Threat model

For the purposes of the current work, we assume an attacker who has physical access to a femtocell base station. The base station uses a landline broadband connection to connect with the mobile operator's core network. This communication channel is encrypted (e.g. using IPSec) and we assume the attacker is not able to wiretap, intercept or modify the communication. The attackers possess one or more mobile devices which enables them to connect to the base station and create signaling traffic such as changing call forwarding settings or sending SMS.

Attacks on the cellular infrastructure can be categorised by three properties: (i) attacks on quality-of-service (QoS), for example by using excessive signaling to affect the performance of the network, (ii) attacks on availability, for example by jamming the radio frequencies, and (iii) attacks on the security and privacy of users. For the cellpot concept, we do not cover attacks on the base station firmware itself, including software-based and hardware-based attacks (e.g. JTAG). However, we note that the security of the firmware is of paramount importance, and current firmware does not offer sufficient protection.

3 Review and Analysis of Attacks in Femtocells

In this chapter we survey various attacks in the femtocell network systems and analyze security issues in femtocell security architecture.

3.1 Analysis of security issues in femtocells

Recent studies demonstrated that femtocells can be easily rooted and turned into rogue base station. In the following section, we discuss why such devices are insecure and outline architecture design issues.

3.1.1 Vulnerabilities of femtocell devices

A well-known problem with femtocells is their system design, which is tailored for minimal costs instead of security. Previous work showed how femtocells can be rooted and how that poses huge risks for both the operators and their customers [11].

Even more concerning is that a single malicious femtocell could poison the whole network, which could then no longer be trusted. We believe that future femtocell hardware will suffer from the same security weaknesses because they will also be tailored for small cost. An analysis of femtocell vulnerabilities shows that they are caused by a combination of three factors: First, the femtocell firmware is built using outdated versions of open source software. Second, it employs a web-based configuration environment, which requires a web server to run on the femtocell. Web servers have a bad security track record and present a broad user-accessible attack vector. Third, the components running on the femtocell are insufficiently isolated from one another. If an attacker succeeds in executing custom code (e.g. through a vulnerability in the web server), she can easily obtain root permissions by rooting the device.

3.1.2 Vulnerabilities of femtocell architectures

It is important to note that traditionally the security of telecommunication networks are based on trust relationships and the fact that it is hard for adversaries to tamper

operator equipment. This has been proven to be problematic in the past and operators face new challenges such as services offered by external gateway providers and massive fraud problems. Again operators are doing the same by essentially introducing a new infrastructure part and relying on a single point of failure which in this case is the HNB. The proliferation of gaining administrative access remotely, open knowledge of the 3GPP standards and related specifications, increasing modern attacking vectors against embedded systems and a risk of giving physical access is making security of not only femtocell devices but also overall infrastructure difficult to control and characterize. Various researchers have evaluated and demonstrated noteworthy attacks from a rogue femtocell and their impacts affecting end-users and mobile operator as well. The presented attacks are irrelevant of a specific operator or system vulnerabilities, instead, they are caused by the vulnerabilities in femtocell security architecture and due to negligence of fundamental 3G security principles. We believe that attacks specifically targeting end users are troublesome for the mobile operators and difficult to mitigate due to nature of the femtocell.

3.2 Attacks against end-users

In this section, we describe attacks against the end-users connected to the compromised femtocell.

3.2.1 Interception attacks

The femtocell acts as a base station serving mobile services to the end-users. If such devices are compromised by exploiting security weaknesses, they can be turn into a rogue base station. These rouge base stations are commonly termed as IMSI-catcher devices. Their main goal is to collect IMSI numbers of mobile devices attached to the base station and to intercept a mobile subscriber's communication (connected to the base station). Authors in [11, 30] demonstrated that if a femtocell is rooted then it can be convert into IMSI-catcher device. In particular, they demonstrated that such compromised devices affects confidentiality of the subscriber data which is one of the important security aspect from operator's view. Their results also claimed that rooted femtocell acts a proxy between the target's phone and the operator. Due to this, such illegal proxy may be difficult to detect. Further, such proxy would give access to target's voice and data communicated over radio interface.

3.2.2 Denial of service attacks

Availability is another important security aspects of mobile communication networks. Rouge femtocells affects this availability aspect via Denial of Service (DoS) attacks. Recent research work demonstrated the possibility of such attacks [11]. They exploited the fact that the *IMSI DETACH MM* message is not authenticated in GSM and 3G networks [22] to launch DoS attacks against the mobile devices connected to rogue femtocell. The attacker can inject these *IMSI DETACH MM* messages via the rooted femtocell to the network and inform the operators that respective users are not available to be paged for incoming mobile services. Once acknowledged by the network, users connected via the rooted femtocell would not receive any mobile terminated services such as voice calls or SMS messages. The important point in this attack is that the non-availability of mobile terminated services is not notified to the users. Their results [11] show that it is possible to perform a large scale DoS attack from rogue femtocells against all subscribers connected via femtocell networks.

3.3 Attacks against the mobile network

In this section, we survey attacks on the Femtocell-enabled cellular network from a compromised femtocell. We describe attacks that do not target mobile phone users directly, but the availability of the network.

3.3.1 Signaling attacks

As discussed earlier, rooted femtocell acts a proxy between the target's phone and the operator. This implies that proxy can be used to inject signaling messages in the core network. As we know that signaling attacks are one of the key threat to the availability of mobile communication networks [9, 27, 31]. If the femtocell is compromised then they can be used to perform signaling attacks on the operator's network. Authors in [11] were able to demonstrate feasibility of such attacks. During their work, they injected malicious traffic to the HNB-GW using attack client and the proxy setup. More importantly to send such malicious traffic, the mobile devices does not need to be connected to the compromised femtocell. Authors claimed that such attack can be automated using their setup to generate signaling flood messages.

3.3.2 Femtocell botnets

If the DoS attacks are carried out from number of compromised femtocells, its impact would be more on the core network components. Authors in [11] investigated this possi-

bility. To make DoS attacks from more femtocells, the attacker needs to remotely control them. Eventually he or she needs to exploit some weaknesses in femtocell devices in order to gain remote root access. The authors were able to discover a remote root access vulnerability to build a botnet of compromised femtocells. They argued about feasibility of making such botnet, provided that it fulfills following conditions:

- During their research, communication between two femtocells were not filtered. In fact, the 3GPP standard mentions that communication between two femtocells is explicitly allowed [3].
- All deployed femtocell devices are same, implies that a vulnerability affecting one devices can be applied on others.
- It would be difficult for the femtocell users to notice that their femtocell device is part of a botnet.
- These devices are always connected to the operator's network via broadband Internet connection.

Due to above conditions, their research demonstrated feasibility of making an army of compromised femtocell devices to perform distributed signaling attacks against the operator's core network. It also claim that it could be possible to evade known detection mechanisms by performing attacks at a low-rate in low-volume as described in [18].

4 Cellpot: A Honeypot Framework for Femtocell Architectures

The cellpot is a novel concept for honeypots inside the core mobile network. Cellpot's purpose is threefold. First, it is a key tool for the network operator to gather intelligence on mobile threats. Second, it acts as a means to protect the core network, and finally, it protects the mobile user. Femtocells are attractive points for honeypot deployment since their stronger signal makes the attacker's mobile device connect to the femtocell instead of the macrocell base station. The cellpot concept consists of three components:

- **Cellpot** A cellpot comprises the original small cell hardware and a custom firmware. Its primary duties are to monitor the signaling traffic and to do anomaly detection. It can also be equipped with means to counter attacks, such as software to rate limit signaling commands, or filters for expensive premium SMS/MMS. In our concept as many cellpots as possible are deployed in order to gain a large coverage and thus increase the chances to catch attacks. The custom firmware has to be certified only once for each type of small cell hardware. Because there are much less types of small cells than there are mobile devices, the costs involved with certification are much smaller for small cell hardware than they are for mobile devices.

- **Peer-to-peer network** Cellpots are interconnected with each other in a peer-to-peer (P2P) network. This network is used to share information between cellpots and to distribute command and control information. The P2P network elects *master nodes* based on the throughput of their landline Internet connection.

- **Honeypot gateway server** The HGS is the central unit of control of all deployed cellpots. It is used by the MNO to centrally collect threat information from the cellpots as well as to issue commands for countermeasures.

To gather intelligence, cellpots interpose between the customer and the core network to detect anomalies in signaling traffic. Cellpots are interconnected with each other in a P2P network. The P2P network has the following duties:

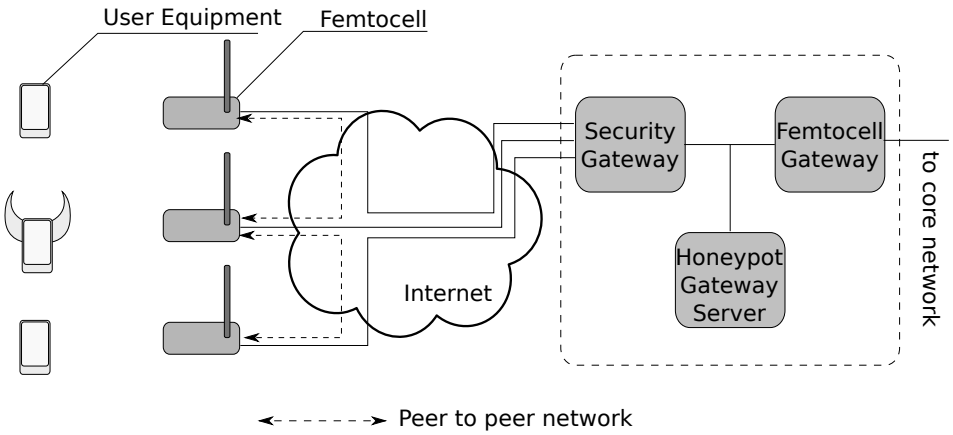

Figure 4.1: Cellpot consists of three components: The cellpot, a peer to peer network and the Honeypot Gateway Server

- **Detect DDoS attacks:** Signalling attacks as shown by Traynor et al. [32] are executed using a large mobile botnet, whose bots do not necessarily share the same location. Because these bots connect to the core network with different base stations, an ongoing attack might seem to be legitimate to a single cellpot. To detect such attacks, cellpots are interconnected with each other with a P2P network and share their information on signaling traffic. If this distributed knowledge indicates an attack, the master nodes will inform the MNO using the HGS.
- **Command and control:** Based on the information received from the cellpot network, the MNO can instruct cellpots to execute countermeasures, e.g. to rate limit or disable certain commands. These commands are sent directly to the master nodes, which distribute them into the P2P network.

We opted to use a P2P network because it significantly increases the scalability of our cellpot infrastructure. This architecture reduces load on the centralized HGS. With this solution, the HGS needs to be connected to a small set of master nodes only. Figure 4.1 illustrates the cellpot architecture.

Cellpot uses *sensors* to record events that could be of interest to collect threat intelligence. A sensor wiretaps the traffic from a communication device and records events of interest. In the case of femtocells there are only two communication interfaces: the radio link and the Ethernet interface. When a sensor detects a suspicious event it can

start to increase the rate with which data is collected. This avoids recording lots of uninteresting events while in the attack case missing important events. In the case of cellpot sensors are also used for threat mitigation. In that case the sensor is acting as a *filter*. For cellpot we envision filters for premium SMS, abnormal signaling traffic and a stolen-devices list.

4.1 Applications of cellpot

In this section we discuss how cellpot can be used by the different stakeholders of the mobile security community. The different stakeholders are mobile network operators, device manufacturers, Computer Emergency Response Team (CERT) organizations, mobile antivirus companies, and academic re- searchers. Honeypots can be categorised by their goals into four types [23]. Honeypots can be used for detection of attacks through e.g. anomaly detection. A prevention honeypot is able to dwarf attacks. Honeypots are used for research to discover patterns and learn about new attacks. To mitigate attacks the intelligence collected by a honeypot can be used to react in a precautionary manner. We believe that cellpot can be categorised in the above four types by interested stakeholders depending on their security requirements. Since the cellpot system is easy for MNOs to integrate into their next generation networks, we discuss new applications. The main advantage for operators to deploy the following applications on the cellpot is to minimize signaling overhead by detecting and preventing various attacks on the small cell itself before it can reach into their core network.

4.1.1 SMS spam prevention

SMS spam is any unwanted text message delivered to mobile users via SMS. This spamming issue continues to grow and constitutes 20-30 of all SMS traffic in Asian markets such as China and India due to the introduction of unlimited text plans [12]. As a consequence of SMS spamming attacks, mobile operators are seeing financial loss due to higher infrastructure and operational costs, poor customer experience, and regulation threats. Typically MNOs deploy various additional solutions within their Signalling System No. 7 (SS7) core network to prevent SMS spam attacks. However such type of solutions introduce additional cost and signaling overhead into the core network. Our cellpot architecture provides a new way for operators to prevent SMS spam. The prevention techniques can be applied directly on small cells and the cellpot gateway. Potential advantage of this method is that operators can detect and block spam messages before they can be sent to the Short Message Service Centre (SMSC), minimizing malicious

SMS related signaling traffic in the core network. The cellpot can be equipped with different filter techniques to block malicious premium rate SMS numbers, mobile malware spreading via SMS messages, and phishing.

4.1.2 Mobile theft prevention

Mobile theft is a rising issue and law enforcement authorities are pushing mobile network operators to tackle it effectively[29]. MNOs deploy Equipment Identity Register (EIR) [13] in their networks and store the identity of stolen or lost phones, typically the International Mobile Equipment Identity (IMEI) number of phones. Operators EIR are automatically connected with other operators to share IMEI database. However deploying additional EIR system introduces additional cost and signaling. Also the system is not effective since attackers usually change the IMEI of the device illegally. The cellpot architecture provides a way to detect stolen mobile phones by uploading IMEI database directly on to the cellpot gateway and small cells. Further mobile data collected in our honeypot system could assist in finding new ways to detect stolen phones despite their IMEI change. This approach does not add SS7 signaling overhead into the core network.

4.1.3 Malware and phishing prevention

The cellpot architecture provides a unique way to monitor mobile data which includes the websites users are trying to connect to. A new anti-phishing framework can be developed using the cellpot architecture similar to Li and Schmitz work in [19]. The cellpot can detect known malicious websites serving malware using services such as MalwareBlacklist.com and inform the operator.

4.2 Legal issues with the cellpot

Our cellpot architecture provides a platform for monitoring mobile traffic including calls, SMS, and data. Depending on its application, the data collected by cellpot can contain user's private information such as International Mobile Subscriber Identity (IMSI) number, call history, and even the browsing history etc. A subset of that data is transferred from small cells to the MNO.

The private nature of this data could raise privacy concerns in some countries. However, we want to stress that the cellpot architecture does not require the MNO to store user's call or SMS data. It is necessary to use certified anonymizing algorithms in the cellpot. Considering that fact that MNOs already provide lawful interception interfaces to government agencies [2], and that they store user's data according to their local laws,

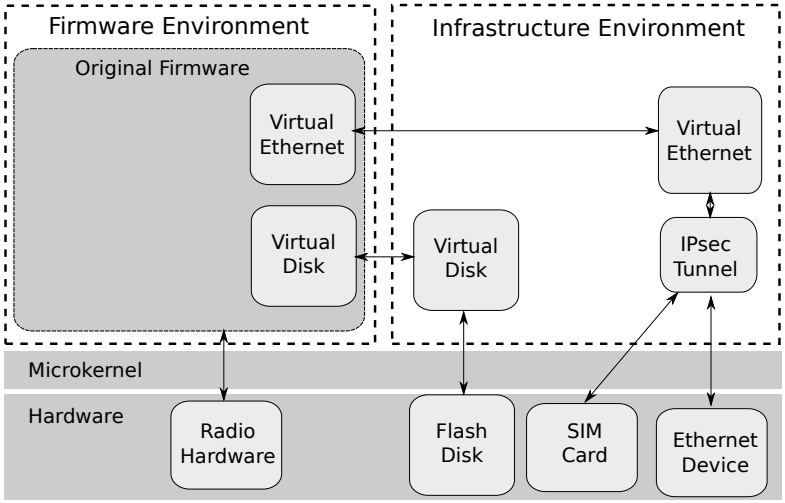

Figure 4.2: The software architecture of a single cellpot consists of two environments; The firmware environment and an infrastructure environment. Only the infrastructure environment is allowed access to cryptographic keys and the Ethernet port. A modern microkernel ensures isolation between the two environments, which are implemented with virtual machines.

we believe that in practice our cellpot will not create legal issues for MNOs during deployment.

4.3 Making cellpot resilient against firmware attacks

As discussed in Sec. 3.1, femtocell devices have security vulnerabilities which may allow an attacker to compromise the femtocell without significant effort, and therefore gain access to the mobile network. Even more concerning is that a single compromised cellpot could poison the whole cellpot P2P network, which could then no longer be trusted. To that end we propose to harden femtocells against rooting by logically partitioning them into two isolated environments. Both environments have separate distinct duties and access distinct pieces of hardware. We call one *firmware environment* (FE), and the second one *infrastructure environment* (IE). This setup is illustrated in Figure 4.2.

The FE has access to the radio hardware and is equipped with a virtual network device. It does not have access to the Ethernet device. We move the entire original firmware

into the FE. The firmware takes care of software defined radio and voice encoding. It uses the virtual network device to communicate with the core cellular network. The FE also hosts the configuration interface. It boots and operates from a virtual disk. The IE in turn has access to the Ethernet port and the flash disk. In particular, its duties are:

- Establishing the link to the core network, using IPSec or similar technology. The key material needed for the link is either hosted directly inside the IE, or in a smart card (e.g. SIM card) that is accessible to the IE exclusively.
- Establishment of a virtual network link to the FE.
- Hosting of the cellpot infrastructure, including its control link and P2P network.
- Establish a virtual disk to host the FE.
- Reset, update, start and stop the FE.

We require the IE to be booted using secure boot. By isolating both environments we assure that rooted firmware can be controlled, e.g. by taking the whole femtocell offline or by resetting the firmware environment. Furthermore the attacker cannot access cryptographic keys or tamper with the cellpot infrastructure. It also solves the problem of costly and time consuming software updates: Certification and validation of the radio stack has to be done only on new firmware versions. Without the costs involved with radio certification and validation, updating and extending the honeypot software can be done frequently.

Contemporary femtocells contain cheap system-on-chip (SoC) components that typically consist of a low power ARM9 core clocked at about 160Mhz and about 64 to 128MB RAM. Currently, these SoCs do not have TrustZone capabilities. Lange et al. showed that virtualization of complex systems like Android is possible on similar embedded systems with the help of a microkernel [17]. Consequently we suggest an implementation using a modern microkernel such as Fiasco.OC as basis, with the partitions being established by virtual machines, similar to the design by Liebergeld et al. [20]. ARM9 and the small amount of memory of current femtocells do not lend themselves to such a system. A system with a Cortex-A9 CPU and about 256MB of RAM enables a performant platform for our software. We argue that the little increase in the total bill of materials is well worth the increase in security.

5 Anomaly Detection of Attacks in Femtocells

This chapter is concerned with the detection of anomalous events in a mobile network that deploys the cellpot concept described in the previous chapter and contains several femtocells. Throughout the chapter we will describe a series of increasing complexity anomaly detection algorithms and test their efficacy for detecting compromised femtocells in a simulated mobile network. In contrast to D4.2 where our target was to detect compromised UEs, in the current deliverable we want to detect compromised femtocells on which several UEs can be connected to. Although the bibliography on anomaly detection is vast, the algorithms we have used cover the most important classes of available methods. The first algorithm we use is *Cumulative Sum* (CUSUM) [23, 4] which is a fast algorithm with a long history in the detection of abrupt changes in several scientific and engineering applications [35],[34]. A more recent method that we describe and evaluate in this chapter is *Local Outlier Factor* (LOF) [7], which detects anomalous measurements exploiting the local deviation of a measurement with respect to its neighbours. Furthermore, a more complex and general algorithm that we experiment with is *Hidden Markov Model* (HMM) [26]. It has been successfully used for time series analysis [21] similarly to our field of application and in several other multidisciplinary domains as a generic machine learning framework [24],[33], [25].

Since all the previous methods require a supervised training phase and are applied on each femtocell separately neglecting the global structure of the network, we also implemented and tested the *Bayesian Robust Principal Component Analysis* (BRPCA) [8]. BRPCA is an unsupervised algorithm that exploits the temporal correlation of measurements across different femtocells of the same network and factorizes a global measurement matrix in a low rank and a sparse component, where each column of these matrices corresponds to a single femtocell. Non zero entries on each column of the sparse component point out anomalous events on the respective femtocell. Similarly to our application domain BRPCA has been successfully applied on a network of traffic sensors [36] for the detection of anomalous conditions in vehicular traffic.

The following sections describe in much greater detail each method along with any algorithmic choices we made with our experimental setup in mind. Moreover, we describe

briefly the simulated mobile network along with the type of measurements that we use. The chapter concludes with the evaluation of the presented algorithms and the findings of our experiments.

5.1 Cumulative sum based detection

The CUSUM test has been proven effective for the detection of DoS attacks [10]. The rationale behind CUSUM is that during an anomalous event, it is expected that several consecutive irregular measurements will arise. CUSUM aggregates these discrepancies in consecutive measurements and if a certain threshold is exceeded signals an anomaly.

Assuming we have N time bins, the basic equation underlying CUSUM is:

$$\begin{aligned} c_j &= c_{j-1} + \max\{0, d_j\} \ , \ j \in \{1, \ldots, N\} \\ c_0 &= 0 \end{aligned} \tag{5.1}$$

where c_j is the anomaly score and d_j is the discrepancy of the measurements in time bin j. To capture the seasonality of the features, where under normal conditions the measurements have a 24 hours period, the discrepancy d_j is defined as,

$$d_j = \sum_{k=1}^{P} \frac{y_{jk} - \bar{y}_{jk}}{\sigma_{jk}} \tag{5.2}$$

where P is the number of different features, $\bar{y}_{jk}$ denotes the mean value of feature type k during the time of the day where time bin j corresponds to, and σ_{jk} its standard deviation. To compute the mean and the standard deviation, a training phase is required where the respective features of a femtocell under normal conditions are recorded.

According to Equation 5.1 there is no resetting mechanism and if an anomaly occurs where the anomaly score c_j is high, it will continue to have high values even if network conditions are restored. To remedy this fact if the discrepancy d_j is low for a number of consecutive measurements we reset c_j to a regular average value. This number is set to 5 for our experiments.

5.2 Local outlier factor based detection

The LOF [7] method measures the outlier-ness of the measurements by examining their sparsity compared to other normal instances. LOF uses all the measurements at a given

time instant simultaneously, therefore we define feature matrix $\mathbf{Y}_i$ for femtocell f_i as

$$\mathbf{Y}_i = \begin{bmatrix} \psi_1 \\ \psi_2 \\ \vdots \\ \psi_N \end{bmatrix} \tag{5.3}$$

where each row $\psi_j \in \mathbb{R}^P$ contains the measurements for every feature during time bin j. Moreover, LOF requires a training set $\mathbf{T}$,

$$\mathbf{T} = \begin{bmatrix} \tau_1 \\ \tau_2 \\ \vdots \\ \tau_L \end{bmatrix}. \tag{5.4}$$

Initially LOF finds the k nearest neighbours of ψ_j and computes its reachability distance from the training set according to:

$$r_k(\psi_j, \tau_m) = \max\{k - distance(\tau_m), d(\psi_j, \tau_m)\}. \tag{5.5}$$

Specifically, the reachability distance of ψ_j from τ_m is in general the true distance of the two, but if they are close it is equal to distance of the k-th nearest neighbour of τ_m. Subsequently the local reachability density of ψ_j is computed as

$$lrd(\psi_j) = \frac{|N_k(\psi_j)|}{\sum_{\tau_m \in N_k(\psi_j)} r_k(\psi_j, \tau_m)} \tag{5.6}$$

where $N_k(\psi_j)$ denotes the k nearest neighbours of ψ_j and $||$ denotes the cardinality of a set. Finally, the anomaly score $LOF(\psi_j)$ is defined as

$$LOF(\psi_j) = \frac{\sum_{\tau_m \in N_k(\psi_j)} \frac{lrd(\tau_m)}{lrd(\psi_j)}}{|N_k(\psi_j)|} = \frac{\sum_{\tau_m \in N_k(\psi_j)} lrd(\tau_m)}{|N_k(\psi_j)| lrd(\psi_j)} \tag{5.7}$$

which is the average local reachability density of the neighbors divided by the local reachability density of ψ_j. A value of approximately 1 indicates that ψ_j is comparable to its neighbours and thus not an outlier. A value below 1 indicates a denser region which would probably be an inlier, while values significantly larger than one indicate outliers.

5.3 Hidden Markov model based detection

HMM is a popular Markovian technique that has been widely applied for anomaly detection [15], [10]. The proposed anomaly detection method is based on an Continuous Density HMM (CDHMM) [26] and consists of two steps: a) the training step, where matrix $\mathbf{T}$ of Equation 5.4 is used to train the HMM, and b) the evaluation step, where the HMM computes the anomaly score for every time instant for femtocell f_i using the measurement matrix $\mathbf{Y}_i$ defined in Equation 5.3.

For the training of the model the Viterbi, the Baum-Welch algorithms [26], [6] and their combination were considered. Based on testing the Baum-Welch algorithm was selected since it performed better. Additionally, in all training cases the distribution of the initial probabilities was set to be uniform. The intuitive explanation for this is that each state corresponds to a specific time of the day and thus it is equally possible for an observation sequence to begin with each state.

To evaluate the anomaly score the trained HMM were used to compute the probability of an observation matrix. This can be written as $P(\mathbf{Y}_i|\lambda)$, where λ are the parameters of the model learned in the training phase. In order to compute this probability the forward part of the forward-backward algorithm is used as outlined in [26],[5].

5.4 Bayesian robust principal component analysis based detection

Another algorithm we implemented to detect anomalous events in a femtocell network is BRPCA [8]. BRPCA is based on the linear dependency of the measurements across different femtocells of the network. The underlying rationale is that the utilization of each femtocell will variate in a similar fashion compared to the others during the course of a day. The biggest advantage of BRPCA compared to the previous algorithms is that it does not require any training as it exploits the global structure of the measurements. The input of BRPCA is a matrix $\mathbf{Y}$ referring to a single feature type for the entire network. Concretely, the input matrix $\mathbf{Y} \in \mathbb{R}^{M \times N}$ is

$$\mathbf{Y} = \begin{bmatrix} \mathbf{y}_1 & \dots & \mathbf{y}_M \end{bmatrix} \tag{5.8}$$

where M is the number of femtocells and N is the number of time bins. Each element corresponds to a single measurement. For example, in our case element y_{ij} is the number of premium SMSs sent by femtocell f_j in time bin i.

Although numerically $\mathbf{Y}$ will have full rank due to small perturbations across different femtocells, in reality its columns are linearly dependent and its actual rank is much

lower. The reason for this attribute is that under normal circumstances the number of premium SMSs sent from the UEs in a femtocell fluctuates similarly for all femtocells in the course of a day. It should be noted that linear dependence models both the case where the number of premium SMSs is approximately constant across different femtocells and also the case where its proportional when femtocells have different capacity. Under this assumption $\mathbf{Y}$ can be written as:

$$\mathbf{Y} = \mathbf{L} + \mathbf{E} \tag{5.9}$$

where $\mathbf{L}$ corresponds to the low rank component of $\mathbf{Y}$ and $\mathbf{E}$ models small magnitude perturbations.

A common method to recover the low rank matrix $\mathbf{L}$ is Principal Component Analysis (PCA) [14], [6] computed using Singular Value Decomposition (SVD) [28]. Initially, by applying SVD, $\mathbf{Y}$ is decomposed as

$$\mathbf{Y} = \mathbf{D}\mathbf{\Lambda}\mathbf{W}^T = \sum_{i=1}^{r} \lambda_i \mathbf{d}_i \mathbf{w}_i^T \tag{5.10}$$

where diagonal matrix $\mathbf{\Lambda} \in \mathbb{R}^{r \times r}$ contains the singular values of $\mathbf{Y}$ sorted in descending order, $\mathbf{D} = \begin{bmatrix}\mathbf{d}_1, \ldots, \mathbf{d}_r\end{bmatrix} \in \mathbb{R}^{M \times r}$ and $\mathbf{W} = \begin{bmatrix}\mathbf{w}_1, \ldots, \mathbf{w}_r\end{bmatrix} \in \mathbb{R}^{r \times N}$ contain the left and right singular vectors respectively and r is the rank of $\mathbf{Y}$. To extract the low rank component $\mathbf{L}$ singular values close to zero and the respective singular vectors are ignored and $\mathbf{L}$ is computed as

$$\mathbf{L} = \sum_{i=1}^{r'} \lambda_i \mathbf{d}_i \mathbf{w}_i^T \tag{5.11}$$

where r' is the number of singular values exceeding a predefined threshold ϵ close to zero and is also equal to the rank of $\mathbf{L}$. If the linear assumption on the observation vectors $\mathbf{y}_1, \ldots, \mathbf{y}_M$ holds it is anticipated that:

$$r' \ll r \leq \min\{M, N\}. \tag{5.12}$$

In a femtocell network traditional PCA will readily recover the low rank matrix $\mathbf{L}$ from the observation matrix $\mathbf{Y}$ in the absence of any anomalous events. However, if a femtocell is attacked and for example starts sending malicious premium SMSs the linearity assumption is violated and PCA collapses since arbitrary magnitude observations can change dramatically the singular values and vectors of the observation matrix even if they correspond to a small percentage of the elements in the observation matrix. On the other hand, in order to detect attacks on femtocells it is crucial to recognize and

isolate any anomalous observations. To accomplish this task we model the observation matrix as

$$\mathbf{Y} = \mathbf{L} + \mathbf{S} + \mathbf{E} \tag{5.13}$$

where as before matrices $\mathbf{L} \in \mathbb{R}^{M \times N}, \mathbf{E} \in \mathbb{R}^{M \times N}$ correspond to the low rank and the noise component, whereas $\mathbf{S} \in \mathbb{R}^{M \times N}$ is a sparse matrix representing anomalous events that cause an arbitrary increase in the observations.

To recover the terms of Equation 5.13 the Bayesian model used by BRPCA is

$$\mathbf{Y} = \underbrace{\mathbf{D}\mathbf{X}\mathbf{\Lambda}\mathbf{W}}_{\mathbf{L}} + \underbrace{\mathbf{B} \circ \mathbf{X}}_{\mathbf{S}} + \mathbf{E}. \tag{5.14}$$

The low rank component $\mathbf{L}$ is modeled as

$$\mathbf{L} = \mathbf{D}(\mathbf{Z}\mathbf{\Lambda})\mathbf{W} \tag{5.15}$$

where $\mathbf{Z}$ is a $(K \times K)$ diagonal matrix with binary entries, $\mathbf{\Lambda} \in \mathbb{R}^{K \times K}$ is a diagonal matrix, $\mathbf{D} \in \mathbb{R}^{M \times K}$, $\mathbf{W} \in \mathbb{R}^{K \times N}$ and the parameter K corresponds to the largest rank that can be inferred for $\mathbf{L}$. The decomposition of Equation 5.15 is similar to SVD, however the matrix corresponding to the singular values is $\mathbf{Z}\mathbf{\Lambda}$ instead of a single matrix. This permits BRPCA to decouple rank learning for singular value learning. The rank of $\mathbf{L}$ is inferred from $\mathbf{Z}$ and is set equal to $\|\mathbf{Z}\|_0$, while the magnitude of the singular values is deduced from $\mathbf{\Lambda}$. The sparse component is factorized as

$$\mathbf{S} = \mathbf{B} \circ \mathbf{X} \tag{5.16}$$

where $\circ$ denotes Hadamard (pointwise) product. Notice again that the proposed model separates the learning of sparseness from the learning of values, such that the zero component in S is exactly zero.

To recover the matrices of Equation 5.14, BRPCA applies a Markov Chain Monte Carlo scheme in order to perform posterior inference. Further details of the algorithm can be found in [8].

5.5 Experimental setup

The mobile network used simulates an area of size 5 km by 5 km, covered by 7 macrocells and 13 femtocells distributed within the area. Macrocells have a large range (> 1 km) and none of them are compromised. The femtocells have a range of 50 m or 20 m and depending on the scenario of each experiment two or none of them are compromised.

The mobile devices move within the boundaries of the simulated area following a random waypoint mobility model, which has been modified to enable UEs to preferentially move to areas covered by femtocells. At the start of the simulation, all UEs are randomly distributed in the area. Each UE then selects a random location in the area and a random speed and proceeds to move to its chosen destination. When it reaches the destination, the UE waits for a random amount of time and then repeats this process. Moreover, each area covered by a femtocell has a rectangular "attraction" area and an attraction probability p_a , i.e. when a UE selects a random destination, it has a probability of p_a of choosing attractor a. The total probability of a UE selecting any of the femtocells in the area is then p_a , where $a \in F$ and F is the set of femtocells in the scenario. If the UE does not select any of the femtocell areas, then it proceeds with the normal destination selection process. Since the "normal" selection process is unaware of attractor areas, the UE may or may not select a femtocell area in this stage. If the UE selects a femtocell area as its destination, then the amount of time it waits is chosen from a different probability distribution specific to that femtocell. Note that as a UE is moving across the simulated area to its destination, it may enter and then exit femtocell areas. Every time the UE attaches to a new macrocell or femtocell, the cell id is recorded in a simulation file. The generation and reception of SMS messages by UEs in the network are also recorded in the same file.

Three snapshots of such a file are illustrated in Figure 5.5. In Figure 5.1(a) the numbers underlined with blue correspond to the unique indices of the UEs, the numbers with green to the respective vector indices, whereas with red are underlined the types of the vectors. Therefore the first line refers to UE 9999 and informs us that vector 79999 contains the cells that it was connected to during the simulation. The contents of vector 79999 are depicted in Figure 5.1(b) where the second column contains a unique event index, the third refers to the time instant the UE connected to a new cell and fourth column refers to the cell index. Similarly, vector 75848 concerns UE 9841 and records the destination addresses that this UE sent an SMS to. The contents of vector 75848 are illustrated in Figure 5.1(c) where the second and third column correspond again to the event index and time instant of the event, whereas the fourth column contains the destination address of the SMS. Using the destination address we can infer whether an SMS is premium or not. The third highlighted line of Figure 5.1(a) informs us that vector 29649 records the incoming SMS to the UE with index 3706.

Since our goal is to detect femtocell attacks, we are not interested on the behaviour of single UEs but instead we have to combine the information of the respective vectors in order to extract features that refer to a femtocell. Initially, we divide the time period of the simulation in constant time bins of 1800 seconds. At each time bin we extract the UEs connected to a specific femtocell and measure the number of incoming and

```
vector 79999  FemtocellNet.ues[9999].mobileNic.sender  cellId:vector  ETV
attr interpolationmode  none
attr source  ueCampedCellId
attr title  "camped cell id, vector"
vector 75848  FemtocellNet.ues[9481].appLayer.udpApps[0]  sentMsgs:vector  ETV
attr interpolationmode  none
attr source  sentMsg
attr title  "msgs sent dest, vector"
vector 29649  FemtocellNet.ues[3706].appLayer.udpApps[0]  recvdMsgs:vector  ETV
attr interpolationmode  none
attr source  recvdMsg
attr title  "msgs recvd src, vector"
```

(a) We highlight with green the measurement vector index, with blue the UE index and with red the measurement type.

79999	87068660	108335.149107027182	106
79999	90295867	111575.149107027182	100
79999	100895009	120729.041409434614	105
79999	108849484	127548.507242309829	104
79999	116666502	134758.289440673907	105
79999	124576426	141993.266949360564	100
79999	124703633	142113.266949360564	106
79999	132165322	179723.782891550142	102
79999	144917657	186323.782891550142	101
79999	149685729	192525.209862422116	106
79999	150594618	193545.209862422116	100
79999	150649460	193605.209862422116	105
79999	150702437	193665.209862422116	116
79999	150855767	193829.673203080142	105
79999	150908227	193889.673203080142	100

(b) Vector 79999 contains a column of event indices (2nd), a column of time instants (3rd) and a column with the cells UE 9999 was connected to.

75848	87477096	108764.854981198907	5510
75848	87589603	108879.854981198907	502
75848	87723327	109017.854981198907	8341
75848	87801203	109096.854981198907	502
75848	87911233	109214.854981198907	2989
75848	103471048	122865.580838445138	8510

(c) Vector 75848 contains a column of event indices (2nd), a column of time instants (3rd) and a column with the destination indices UE 9481 sent SMSs to. Premium SMSs are identified by the destination index.

Figure 5.1: Format of a simulation file. There are three different kind of vectors concerning the cells a UE was connected to, the SMSs it sent and the SMSs it received.

outgoing SMSs as well as the number of premium SMSs sent. Therefore, assuming that the network contains M femtocells $\{f_i\}_{i=1...M}$ and the simulation extends N time bins, feature matrix $\mathbf{Y}_i$ arises for femtocell f_i,

$$\mathbf{Y}_i = \begin{bmatrix} \mathbf{y}_1^i & \mathbf{y}_2^i & \mathbf{y}_3^i \end{bmatrix} \tag{5.17}$$

where $\mathbf{Y}_i$ is an $(N \times 3)$ matrix. Columns $\mathbf{y}_1^i$, $\mathbf{y}_2^i$ and $\mathbf{y}_3^i$ correspond to the number of incoming, outgoing and premium SMSs in f_i during the simulation. These feature matrices $\{Y_1, \ldots, Y_M\}$ are the input to the LOF, CUSUM and HMM methods. The BRPCA uses a single feature type (number of premium SMS) and the input of the algorithm is the feature matrix of Equation 5.8.

5.6 Experimental evaluation on various attack scenarios

The efficiency of the presented algorithms is demonstrated under several different attack scenarios that cover a wide range of attacks that can occur in a real network . In the first series of experiments the UEs are active 24 hours a day while in the second series the UEs have a diurnal cycle and are active for between 14 to 16 hours per day. When the UEs are operating according to a diurnal pattern, the UEs are immobile during their inactive period. Each series of experiments consists of five different scenarios where different values for two parameters of the simulated network are considered.

The first parameter denoted as $premProb$ is the probability that a UE sends a premium SMS. The second parameter concerns the attack type and can have three different values. The first value is "no attack". The second one is "periodic" where the compromised femtocells send premium SMS periodically. In particular a UE sends 1 premium SMS every hour as long as it is connected to a compromised femtocell. The third attack type is "bursty" where a compromised femtocell generates and sends a burst of premium SMSs (1 up to 4), once at/near attach time of a UE.

In the experiments where an attack occurs femtocells f_3 andf_{12} are compromised so anomalies should be detected on them. For each simulation scenario we provide two different types of diagrams for each method. The first type depicts the anomaly score for each femtocell at each time bin using a jet colormap. This colomap will use the range of the measurements and assign blue to the lowest and red to the highest. However, one should keep in mind that red entries can correspond to small score in terms of absolute value. In case of an attack we expect many red entries concentrated on the columns corresponding to femtocells f_3 and f_{12}. For BRPCA the anomaly score is equal to the absolute value of matrix $\mathbf{S}$ in Equation 5.13. The second type of diagrams demonstrate the accumulated anomaly score for each femtocell per day. No attack

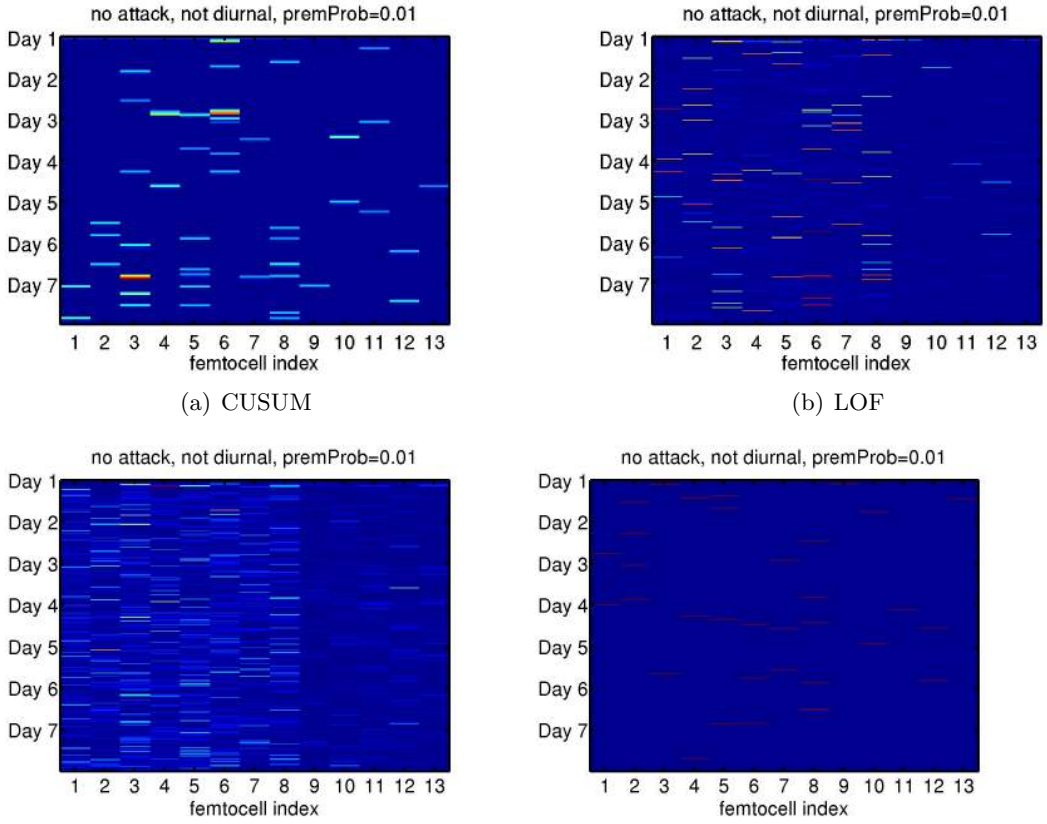

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.2: Anomaly scores for all femtocells throughout the simulation for the CUSUM, LOF, HMM and BRPCA algorithms respectively when no attack occurs, $premProb = 0.01$ and devices are active 24 hours a day. Some sporadic fluctuations appear. BRPCA is more stable compared to the others.

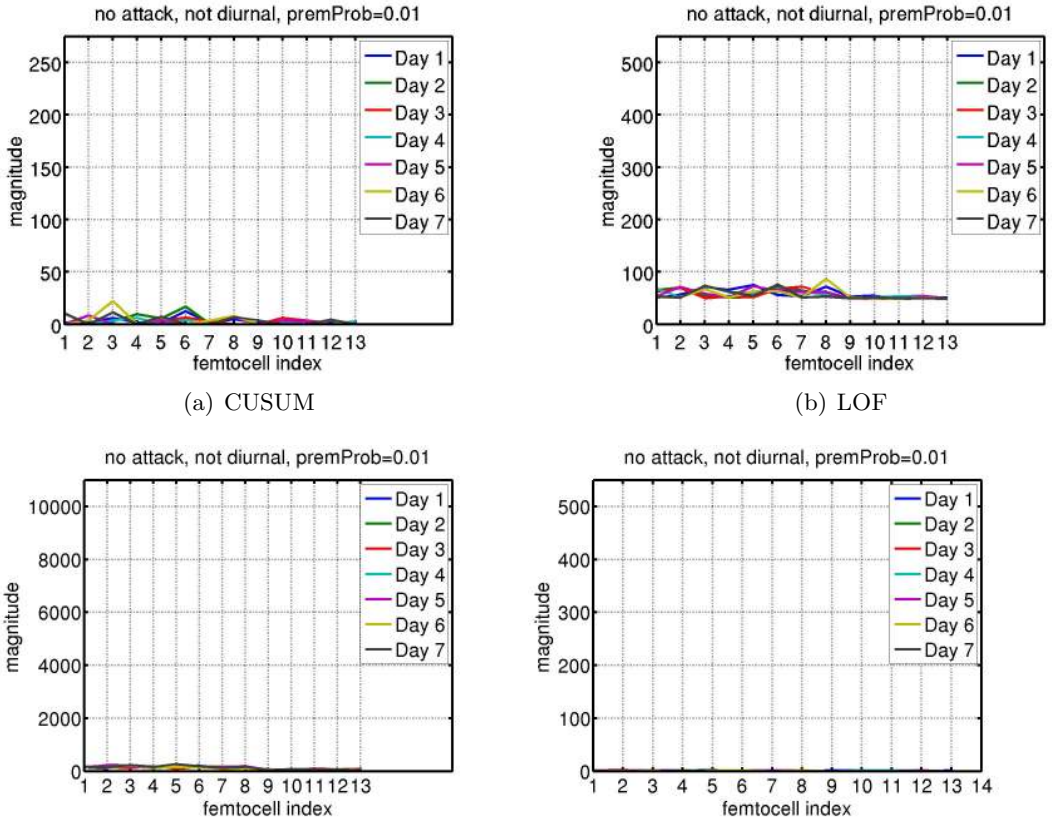

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.3: Aggregated per day anomaly scores of the femtocells for the CUSUM, LOF, HMM and BRPCA algorithms respectively when no attack occurs, $premProb = 0.01$ and devices are active 24 hours a day. Notice the small discrepancies of CUSUM and LOF. The respective score of BRPCA is close to zero.

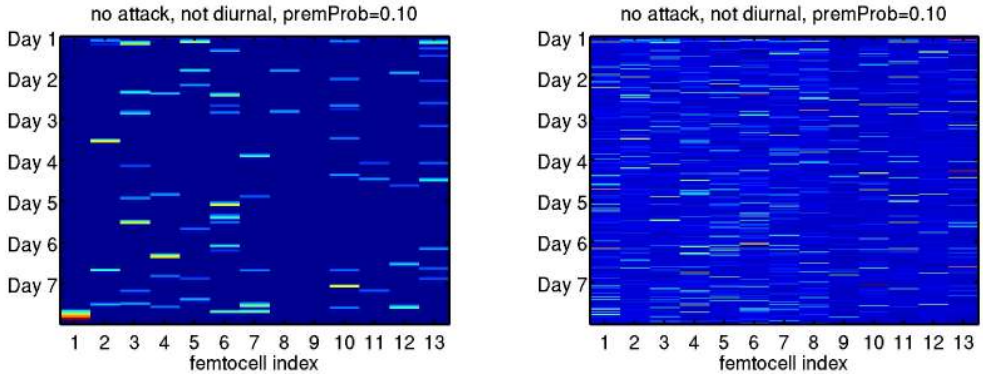

(a) CUSUM

(b) LOF

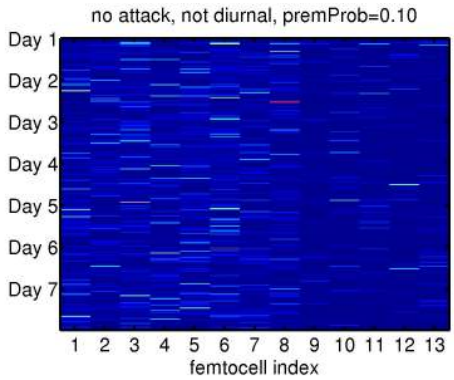

(c) HMM

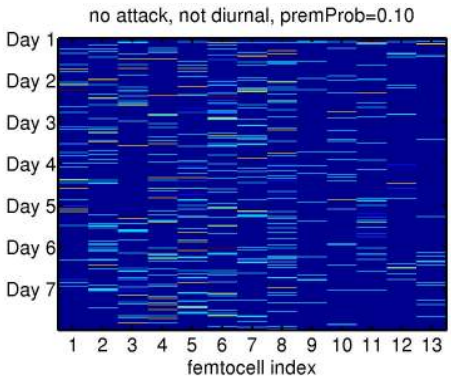

(d) BRPCA

Figure 5.4: Anomaly scores for all femtocells throughout the simulation for the CUSUM, LOF, HMM and BRPCA algorithms respectively when no attack occurs, $premProb = 0.1$ and devices are active 24 hours a day. As $premProb$ has increased, the anomaly score fluctuates more heavily around zero.

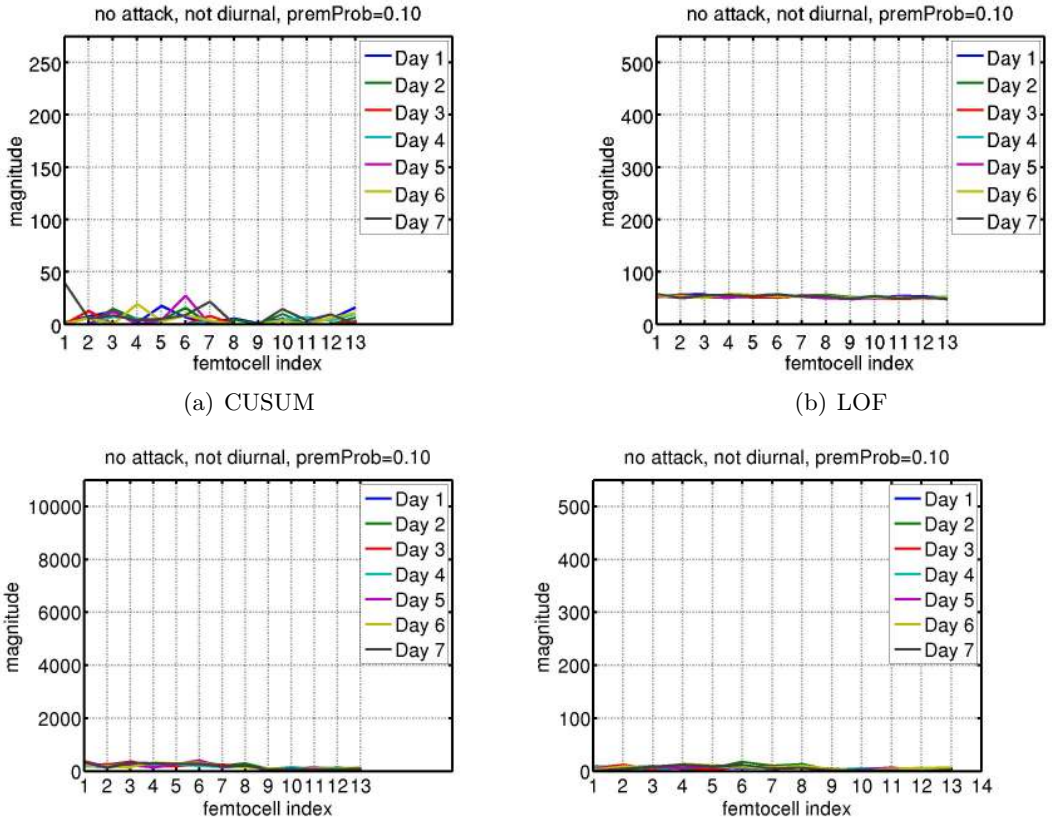

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.5: Aggregated per day anomaly scores of the femtocells for the CUSUM, LOF, HMM and BRPCA algorithms respectively when no attack occurs, $premProb = 0.1$ and devices are active 24 hours a day. Notice the significant increase of the error score for CUSUM. However, by setting a relatively high threshold no false anomalies will be detected. LOF and BRPCA are still stable.

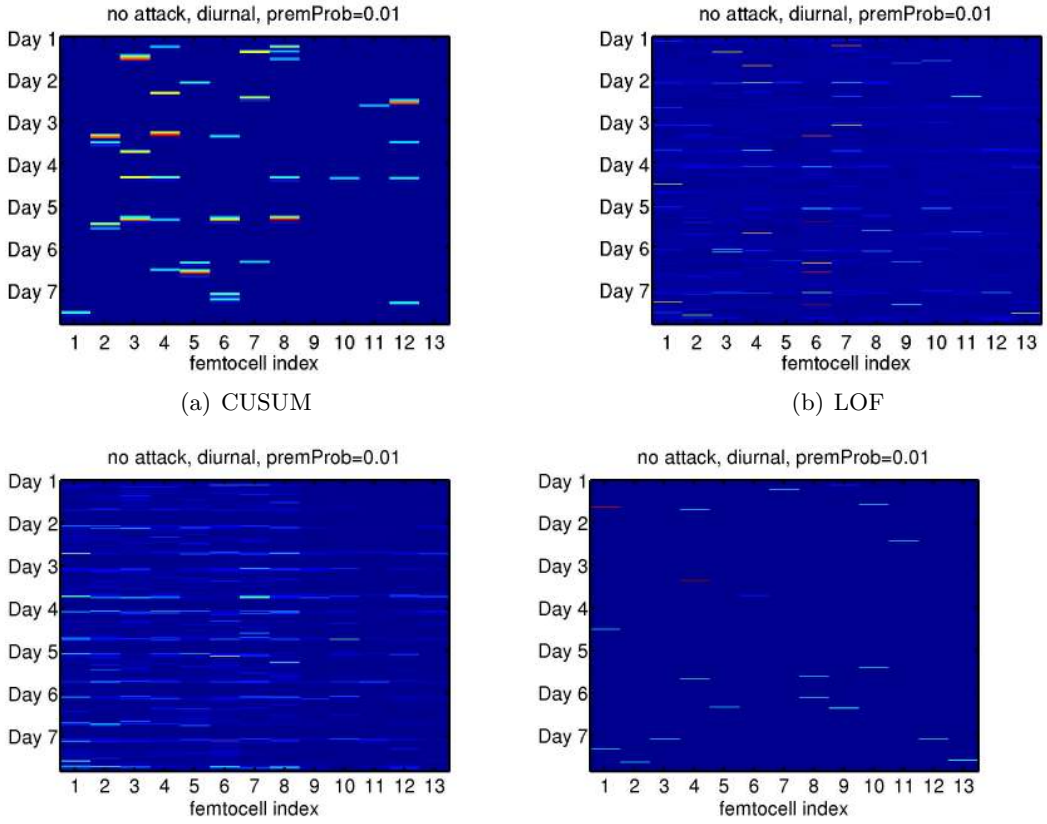

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.6: Anomaly scores for all femtocells throughout the simulation for the CUSUM, LOF, HMM and BRPCA algorithms respectively when no attack occurs, $premProb = 0.01$ and devices have a diurnal cycle. More entries are zero with some minor exceptions that are still close to zero.

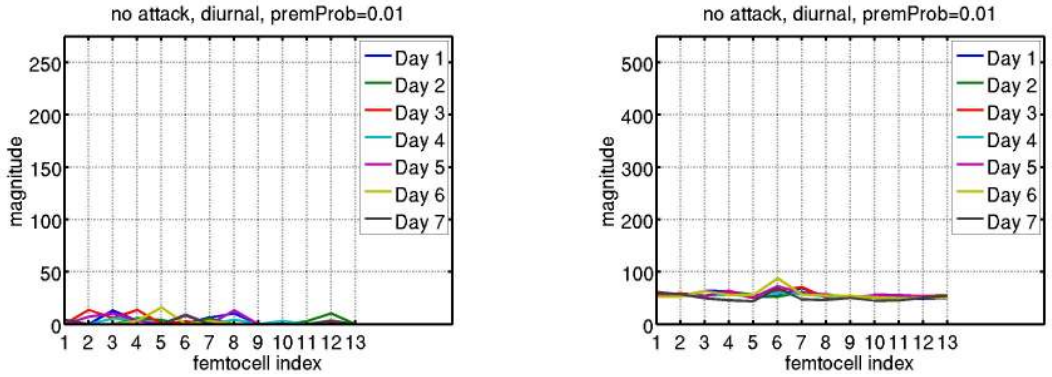

(a) CUSUM

(b) LOF

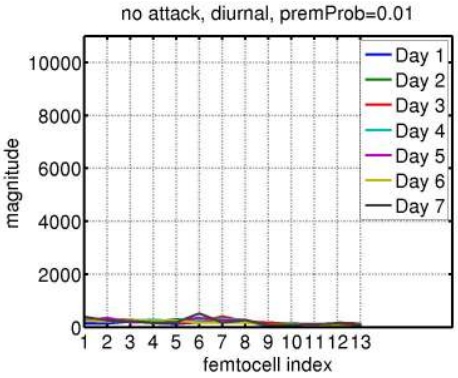

(c) HMM

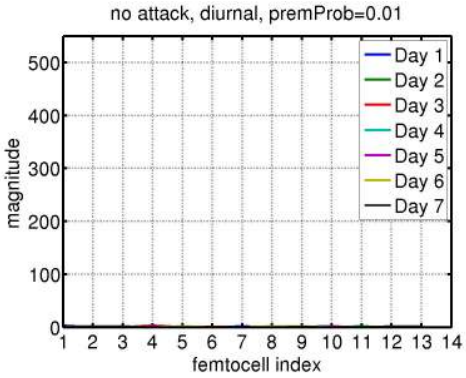

(d) BRPCA

Figure 5.7: Aggregated per day anomaly scores of the femtocells for the CUSUM, LOF, HMM and BRPCA algorithms respectively when no attack occurs, $premProb = 0.01$ and devices have a diurnal cycle.

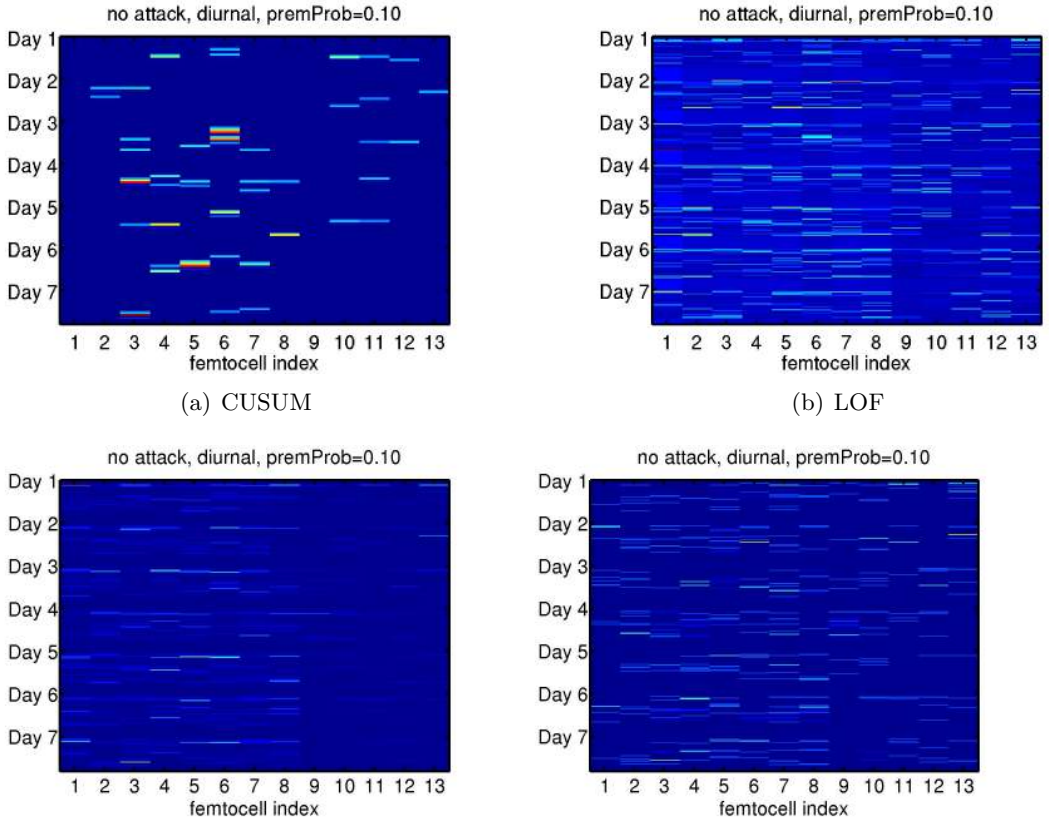

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.8: Anomaly scores for all femtocells throughout the simulation for the CUSUM, LOF, HMM and BRPCA algorithms respectively when no attack occurs, $premProb = 0.1$ and devices have a diurnal cycle. Notice that the anomaly score for BRPCA is practically zero compared to the other methods, where some minor fluctuations occur.

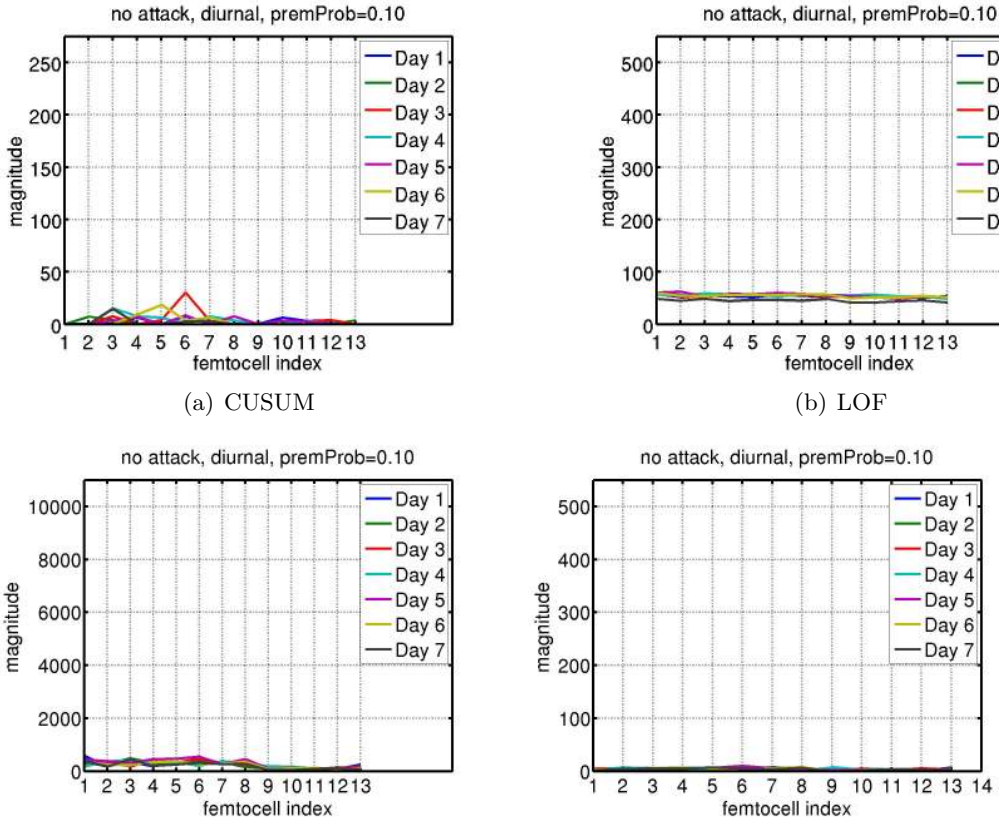

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.9: Aggregated per day anomaly scores of the femtocells for the CUSUM, LOF, HMM and BRPCA algorithms respectively when no attack occurs, $premProb = 0.1$ and devices have a diurnal cycle. As the probability of premium SMSs increased, more non zero anomaly scores arise.

scenarios are illustrated in Figures 5.2, 5.3, 5.4, 5.5, 5.6, 5.7, 5.8, 5.9. The stability of the BRPCA method is of special interest although no spurious alarms are raised by any of the other methods. This stability is attributed on the global view that BRPCA has on the data. CUSUM is the most unstable because small fluctuations on the measurements get accumulated over time, while both LOF and HMM perform equally well.

Cases of "bursty" attacks are shown in Figures 5.10, 5.11, 5.12 and 5.13 where all methods can identify that femtocells f_3, f_{12} are compromised. Among the evaluated methods BRPCA and HMM are the most stable. It should also be noted that although CUSUM detects the compromised femtocells it does not have a constant anomaly score for the anomalous measurements. Its anomaly score is increasing over time as discrepancies are accumulated. This is evident in the jet colomap plots where the color of the columns of the compromised femtocells varies smoothly from blue to red over time. Moreover, when femtocells do not have a diurnal cycle the anomaly curves per day for CUSUM are different since the anomaly score of Equation 5.1 is never reset. The HMM algorithm performs well and the anomaly scores for the compromised femtocells have very high values.

Scenarios of periodic attacks are illustrated in Figures 5.14, 5.15, 5.16, 5.17, 5.18, 5.19, 5.20 and 5.21. Again all methods perform well, however LOF has some small instabilities in the case of Figures 5.14 and5.15. It is also worth commenting on the depiction of the diurnal cycle of the attacks in the corresponding figures. The anomaly matrices have on the columns of the compromised femtocells low values depicted as blue entries during the time of inactivity, followed by red entries when devices are active. All methods perform well and compromised femtocells can easily be detected. We should point out that the blue stripes in the case of HMM that are shown in Figures 5.14, 5.16, 5.18 and 5.20 for the compromised femtocell f_{12} are due to very big anomaly scores that affect the scaling of the colormap. The HMM algorithm works very well as is shown in the per day diagrams.

The methods we presented in this chapter are indicative of broader families of anomaly detection algorithms, ranging from simple control charts to more sophisticated machine learning methods. In general the tested algorithms can accurately detect compromised femtocells. However CUSUM is not as accurate and stable as the others. On the other hand particularly encouraging is the performance of BRPCA which is the most stable method and in addition does not require any training. Moreover, the HMM algorithm works very well but requires a cumbersome training phase and more than twice the computational time which can be a significant drawback for larger networks.

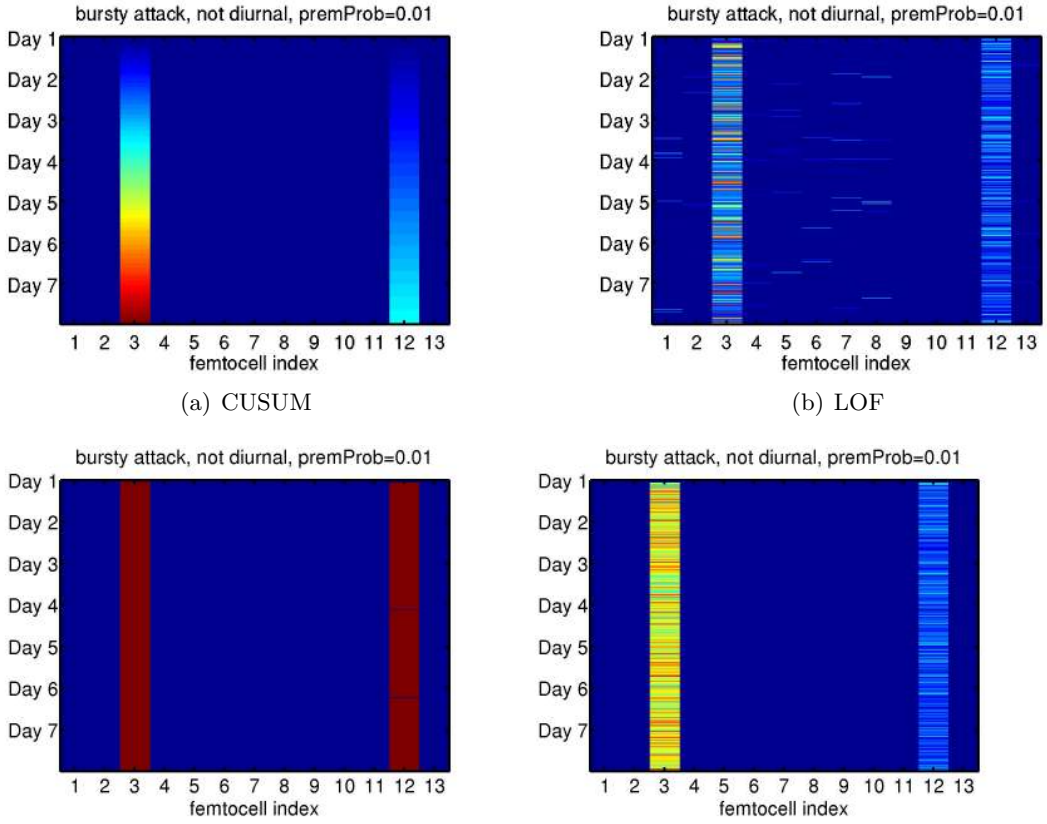

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.10: Anomaly scores for all femtocells throughout the simulation for the CUSUM, LOF, HMM and BRPCA algorithms respectively under a bursty attack, with $premProb = 0.01$. The devices are active 24 hours a day. Notice the high values on columns 3 and 12 corresponding to the compromised femtocells. In the case of CUSUM the color in these columns varies smoothly from blue to red as discrepancies are constantly accumulated.

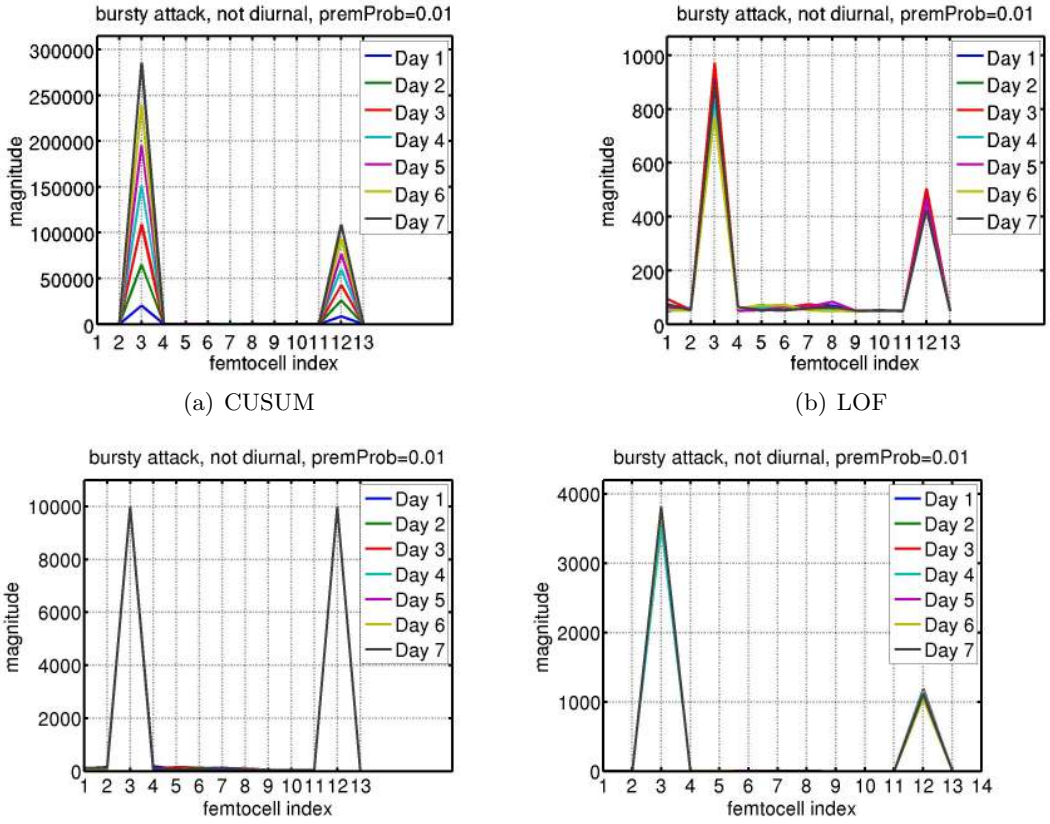

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.11: Aggregated per day anomaly scores of the femtocells for the CUSUM, LOF, HMM and BRPCA algorithms respectively under a bursty attack with $premProb = 0.01$. The devices are active 24 hours a day. The curves for BRPCA are nearly identical for all days. The anomaly curves for CUSUM have higher values every day as discrepancies accumulate over time for the compromised femtocells.

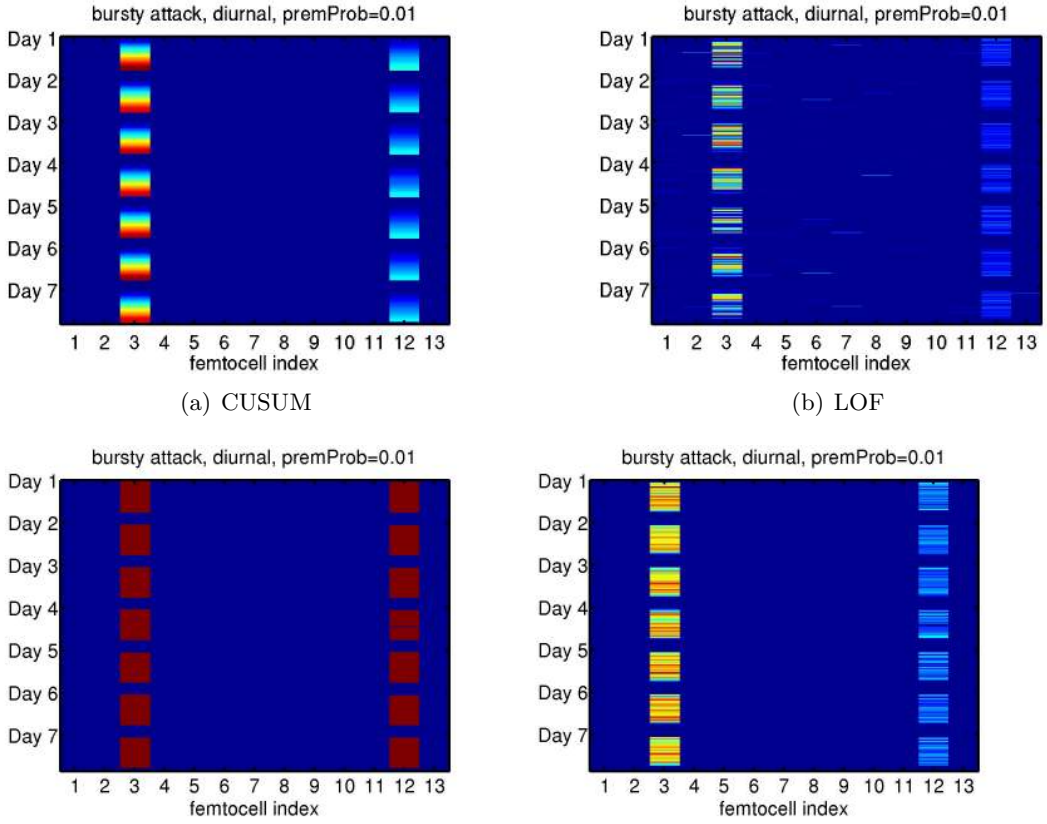

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.12: Anomaly scores for all femtocells throughout the simulation for the CUSUM, LOF, HMM and BRPCA algorithms respectively under a bursty attack with $premProb = 0.01$. The devices have a diurnal cycle and this is depicted in the diagram where only periodic blocks of the columns corresponding to compromised femtocells have a high value. Notice again the stability of BRPCA.

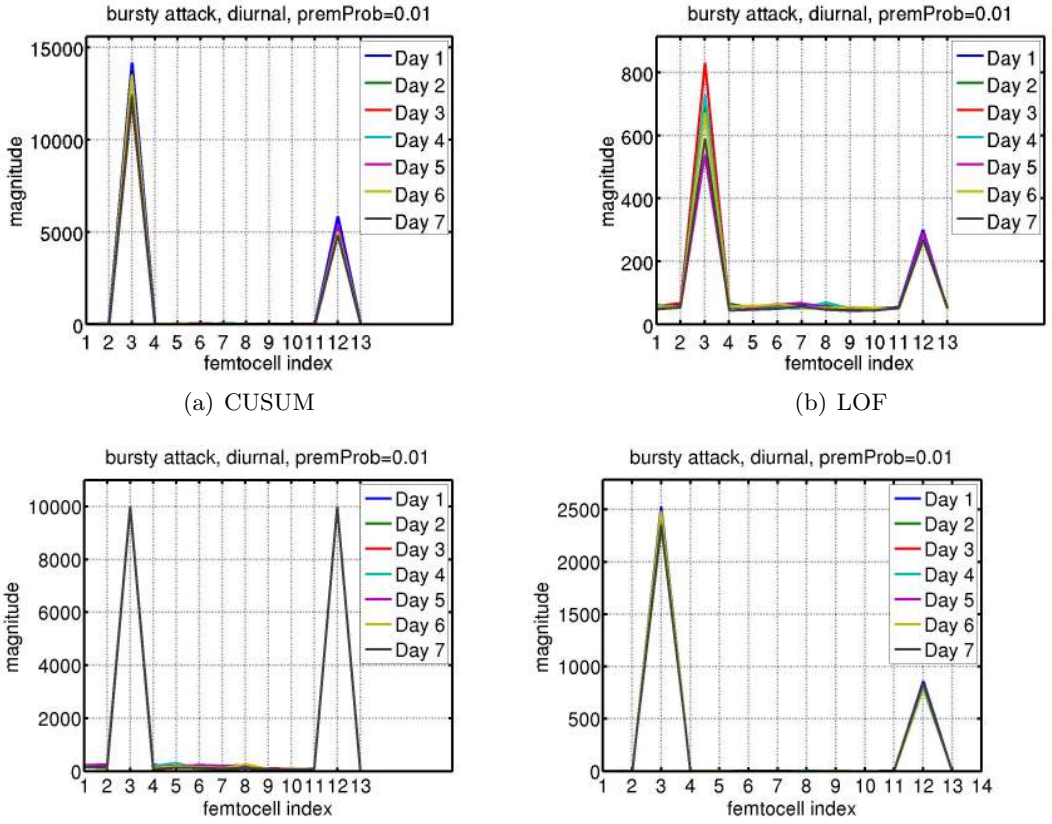

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.13: Aggregated per day anomaly scores of the femtocells for the CUSUM, LOF, HMM and BRPCA algorithms respectively under a bursty attack with $premProb = 0.01$. The devices have a diurnal cycle and all methods perform very well.

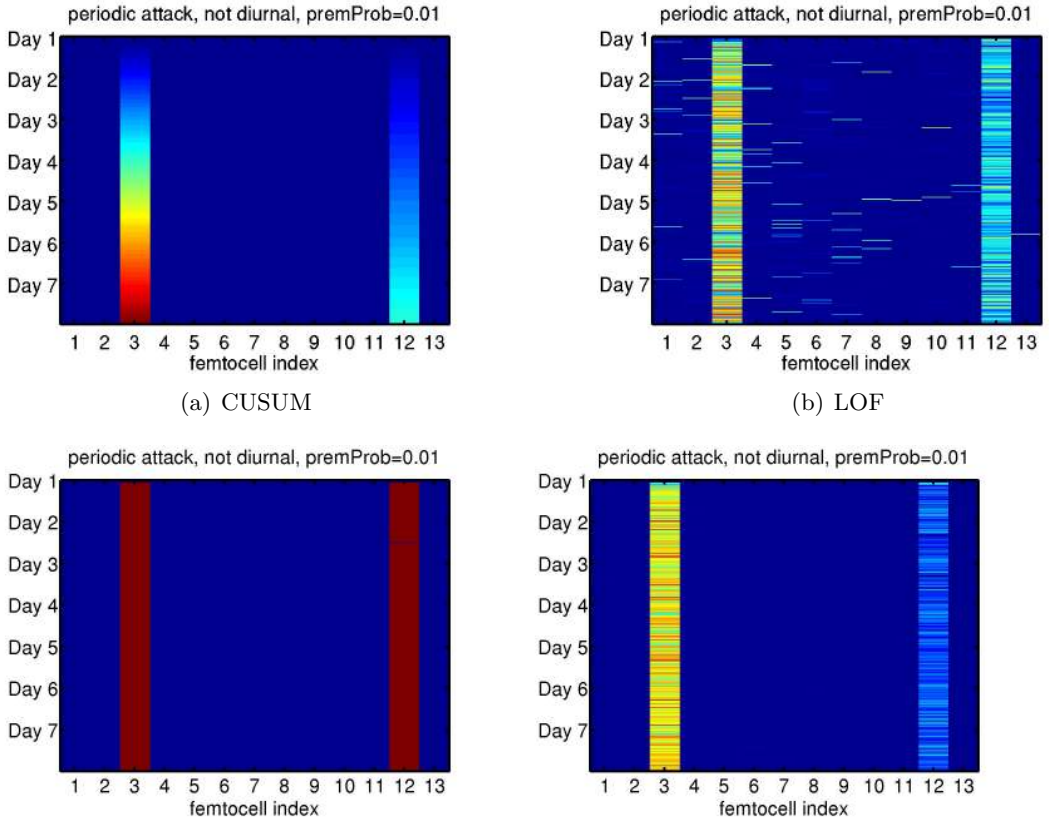

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.14: Anomaly scores for all femtocells throughout the simulation for the CUSUM, LOF, HMM and BRPCA algorithms respectively when a periodic attack occurs, $premProb = 0.01$ and devices are active 24 hours a day. Notice the sporadic instabilities of LOF.

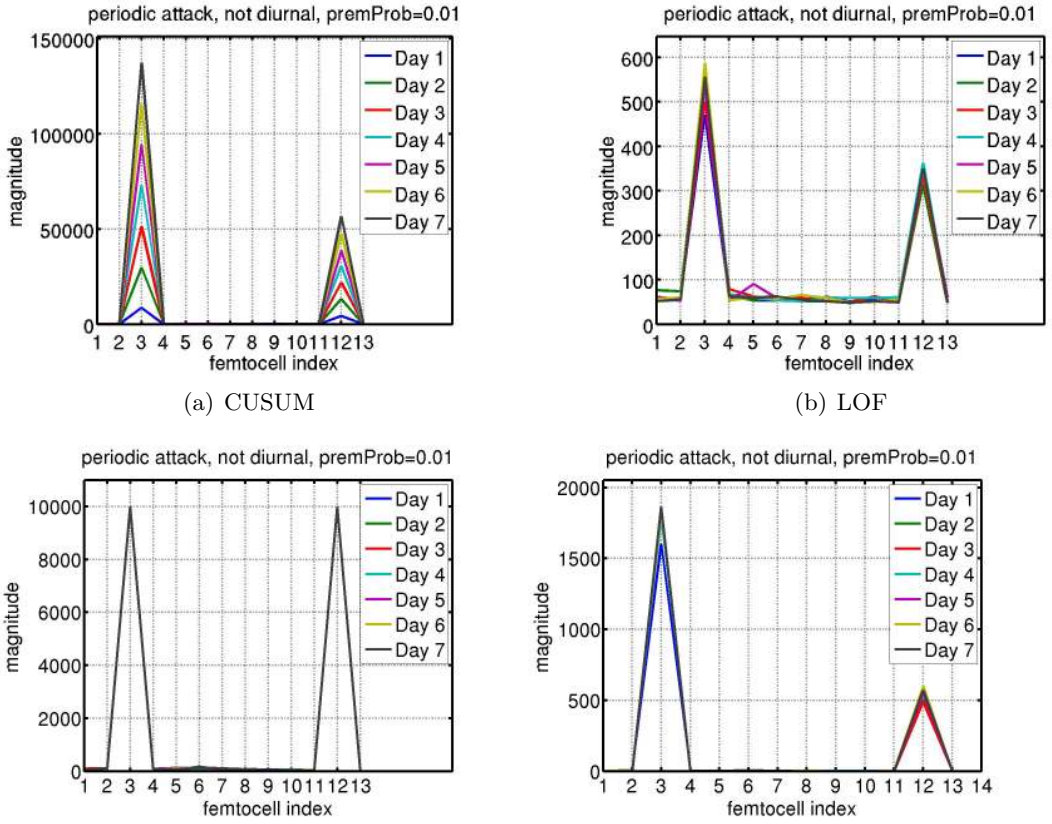

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.15: Aggregated per day anomaly scores of the femtocells for the CUSUM, LOF, HMM and BRPCA algorithms respectively when a periodic attack occurs, $premProb = 0.01$ and devices are active 24 hours a day. CUSUM curve get higher every day as no resetting of the anomaly score occurs. Also LOF has some small instabilities. BRPCA continues to be pretty robust.

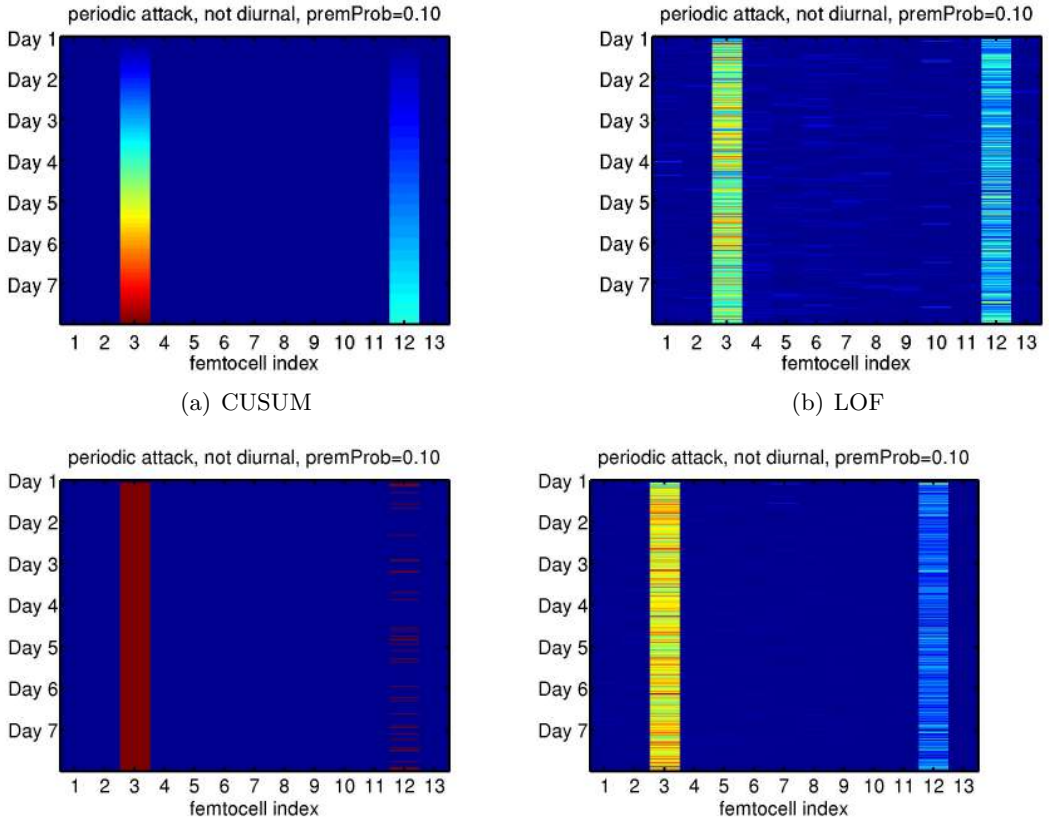

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.16: Anomaly scores for all femtocells throughout the simulation for the CUSUM, LOF, HMM and BRPCA algorithms respectively when a periodic attack occurs, $premProb = 0.1$ and devices are active 24 hours a day. Notice some sporadic non zero entries for LOF in uncompromised femtocells.

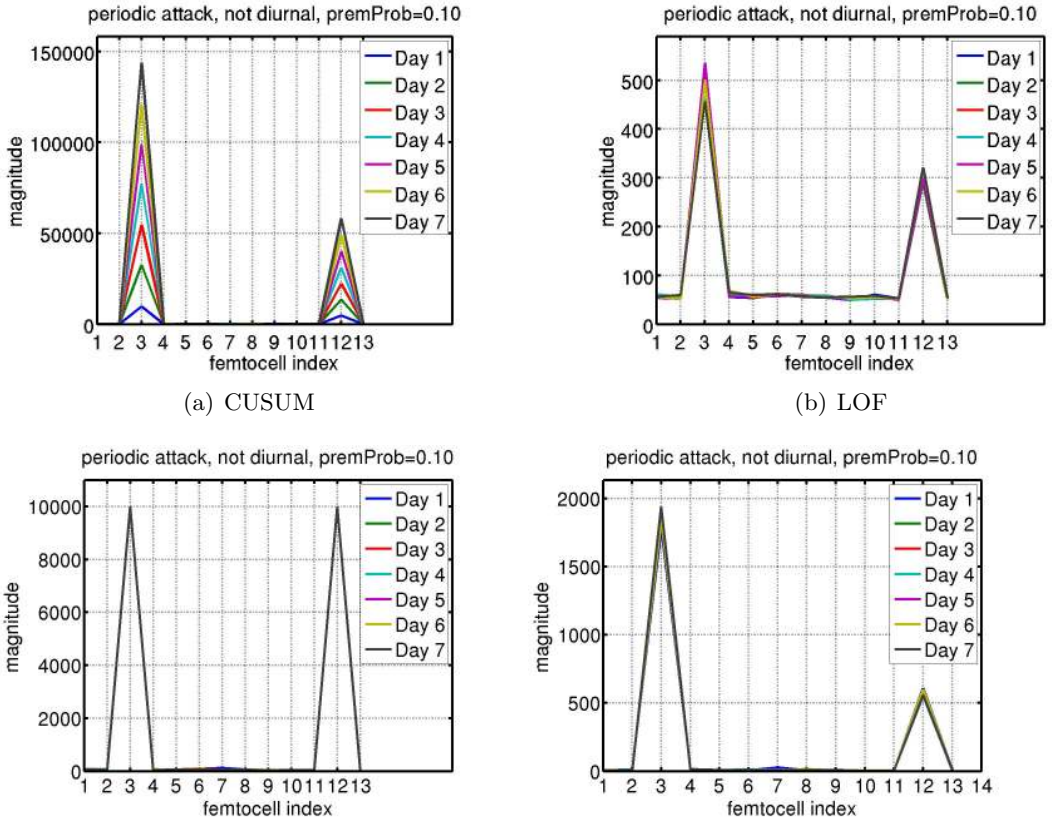

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.17: Aggregated per day anomaly scores of the femtocells for the CUSUM, LOF, HMM and BRPCA algorithms respectively when a periodic attack occurs, $premProb = 0.1$ and devices are active 24 hours a day. All methods perform very well.

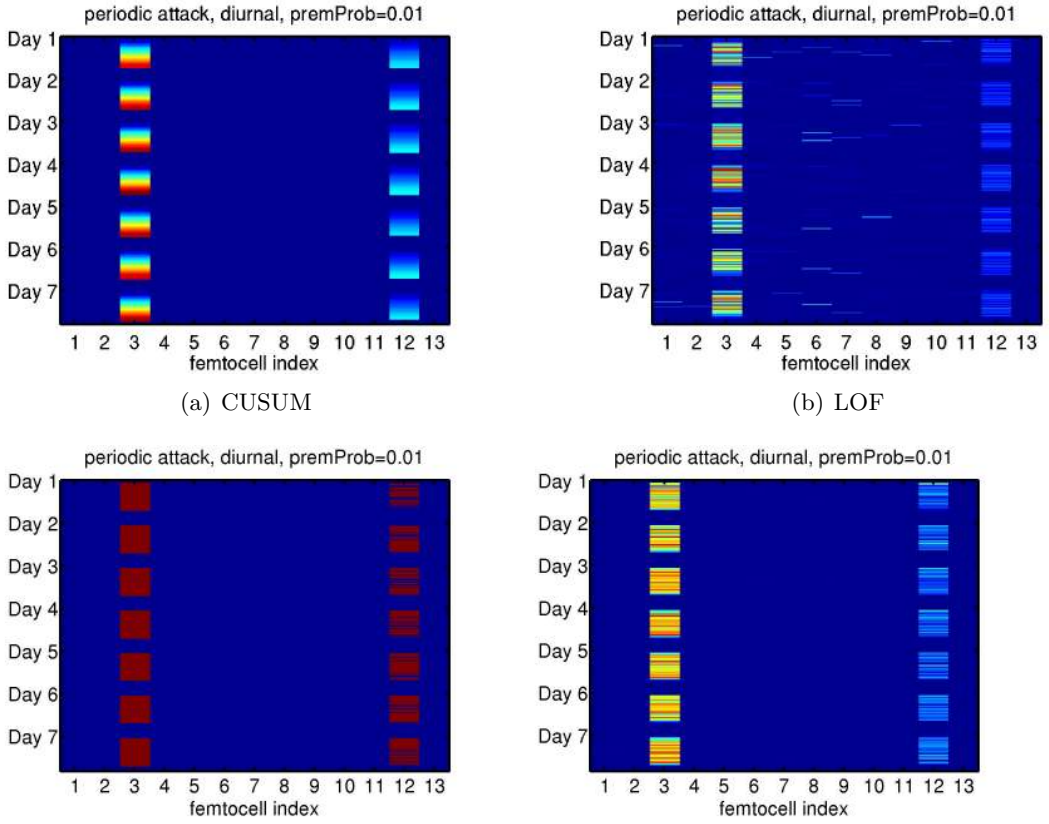

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.18: Anomaly scores for all femtocells throughout the simulation for the CUSUM, LOF, HMM and BRPCA algorithms respectively when a periodic attack occurs, $premProb = 0.01$ and devices have a diurnal cycle. Notice the depiction of the diurnal pattern on the diagram.

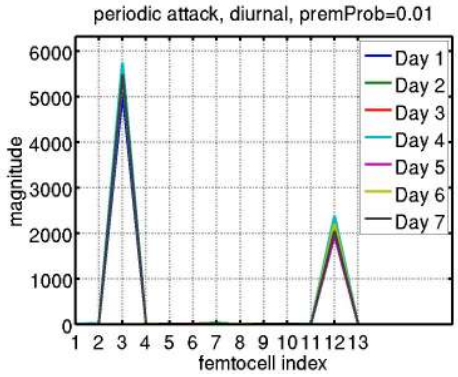

(a) CUSUM

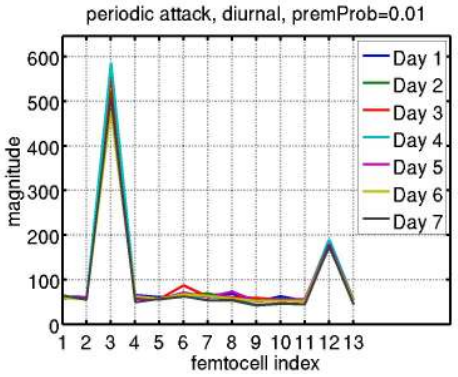

(b) LOF

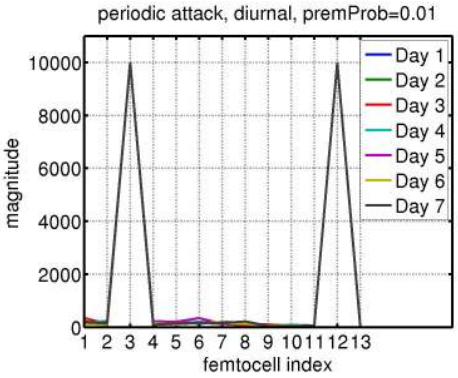

(c) HMM

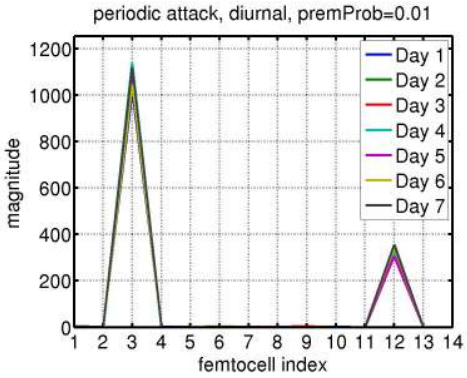

(d) BRPCA

Figure 5.19: Aggregated per day anomaly scores of the femtocells for the CUSUM, LOF, HMM and BRPCA algorithms respectively when a periodic attack occurs, $premProb = 0.01$ and devices have a diurnal cycle. Notice the minor instabilities of LOF.

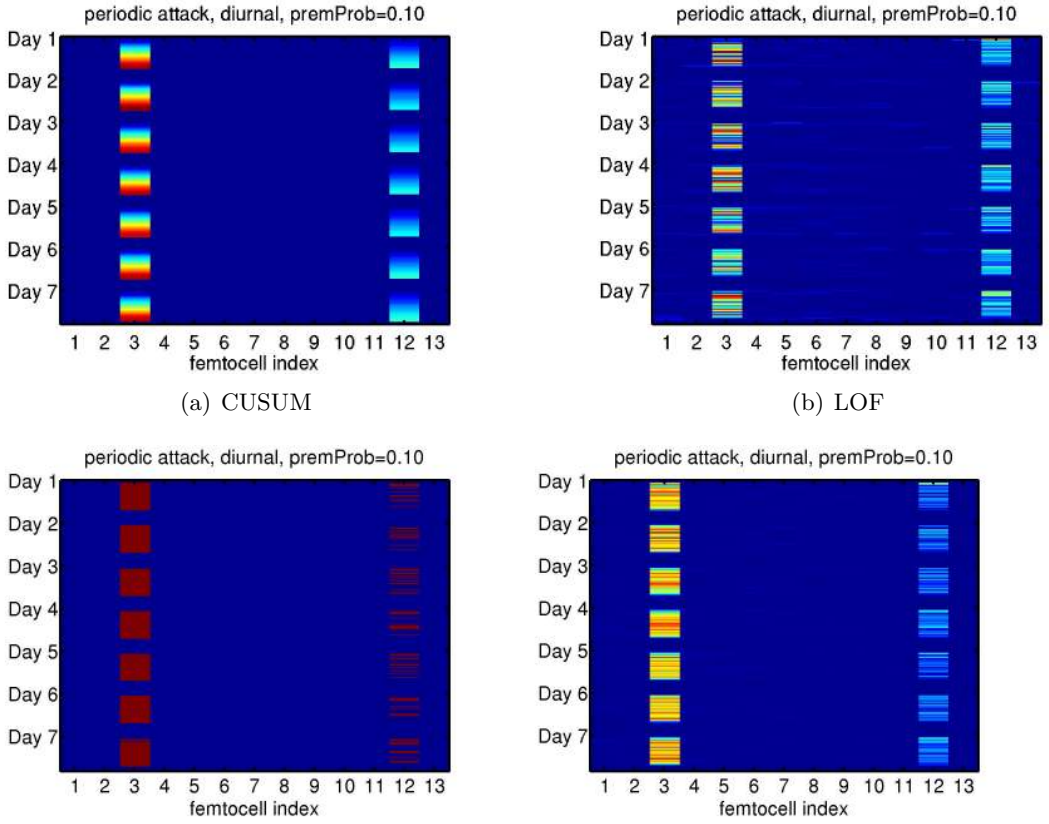

(a) CUSUM

(b) LOF

(c) HMM

(d) BRPCA

Figure 5.20: Anomaly scores for all femtocells throughout the simulation for the CUSUM, LOF, HMM and BRPCA algorithms respectively when a periodic attack occurs, $premProb = 0.1$. The devices have a diurnal cycle. Again the diurnal pattern is depicted.

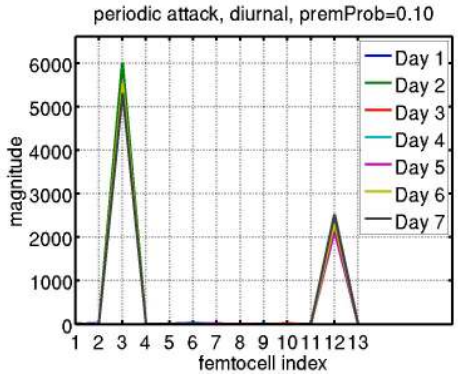

(a) CUSUM

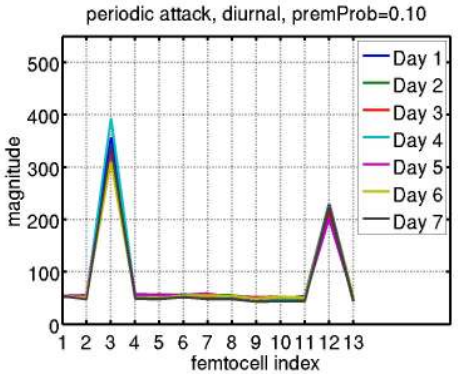

(b) LOF

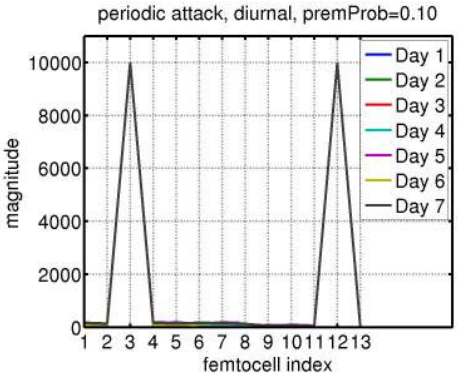

(c) HMM

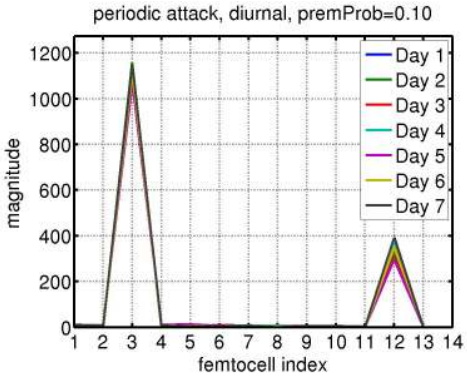

(d) BRPCA

Figure 5.21: Aggregated per day anomaly scores of the femtocells for the CUSUM, LOF, HMM and BRPCA algorithms respectively when a periodic attack occurs, $premProb = 0.1$ and devices have a diurnal cycle. All methods perform well.

6 Conclusions

In this document, we presented a review of the security vulnerabilities of femtocells and discussed the femtocell-specific attack vectors due to potential exploitation of these vulnerabilities. Femtocell devices are critically situated for an attacker since they have easy physical access to the device, which allows them to physically attack and compromise the device. In addition, femtocell devices typically run on outdated open source software, which make them vulnerable to exploitation attacks. Femtocell devices also run a web server to allow easy configuration of the device, which introduces another attack vector that can be exploited by an adversary in order to compromise the device due to insufficient isolation of the device components. Furthermore, due to the vulnerabilities of the femtocell security architecture, a compromised femtocell device allows (almost) direct access to the end-users and the mobile services of the operator, making femtocells an attractive target.

Once a femtocell is compromised, the attacker can launch attacks against the end-user, such as interception and DoS attacks (Sec. 3.2), or attacks against the mobile network, including signaling attacks (Sec. 3.3). In order to monitor femtocells for anomalies which may arise due to compromised devices and other factors, and to enable anomaly detection and mitigation against identified attacks, we propose the concept of the cellpot (Chp. 4), which is comprised of femtocell-based honeypots connected via a peer-to-peer network that allows the distributed sharing of information regarding anomalies among the cellpots, and enabling distributed anomaly detection and control for defending against identified attacks. In order to improve the security of the femtocell device against attacks exploiting vulnerabilities in the firmware, we propose a software architecture for the femtocell that protects the core network even if the femtocell firmware is compromised. This software architecture uses virtualization in order to separate the femtocell device into two environments: the firmware and the infrastructure environments. This separation allows us to protect the core network since only the infrastructure environment has access to the cryptographic keys used for authentication and authorization, and to the Ethernet port used to connect the device to the mobile network.

In order to fully realize the monitoring and protection of femtocells via cellpots, we also investigated anomaly detection methods that will take as input the data collected by the cellpots, e.g. via the honeypot gateway server, and analyze the femtocell data in

order to identify any anomalies occurring in the femtocell architecture. In future stages of this work, once an anomaly is identified and classified, defensive actions can be taken via the cellpot in order to mitigate against the anomaly or attack. For example, if a femtocell is identified as compromised, then mobile traffic to/from the femtocell can be further scrutinized and filtered based on firewall-like rules in order to protect the core network and the mobile users.

The anomaly detection methods we considered were based on the following: cumulative sum (CUSUM), local outlier factor (LOF), hidden Markov models (HMM), and Bayesian robust principal component analysis (BRPCA), which were chosen due to their various advantages. For example, CUSUM is a computationally efficient and fast detection method, while HMM-based anomaly detection methods have been shown to provide good detection performance after proper supervised learning. We evaluate these methods using simulation data from mobile networks with compromised femtocells. In the evaluation scenario, we consider currently the most common malware type, premium SMS messages, and apply the anomaly detection methods in order to identify the anomaly and its source, i.e. which of the femtocells are compromised. Our results show that all of the considered methods are able to detect the anomaly. However, we see that despite its efficiency, CUSUM is the least stable and least accurate, while both HMM and BRPCA are the most stable and have consistently good detection performance. We recommend the use of BRPCA over HMM since BRPCA does not require supervised training and is computationally more than two times faster than HMM, which makes BRPCA an attractive candidate for anomaly detection in large scale mobile networks with femtocells.

Future work in the security and protection of femtocells would need to consider appropriate mitigation methods and how they can be integrated and realized with the proposed cellpot concept. In addition, more in-depth performance analysis of anomaly detection algorithms as implemented within the cellpot framework would be desirable in order to evaluate whether they can be used in (near) real-time for anomaly detection within femtocells.

Bibliography

[1] 3GPP. Security of Home Node B (HNB) / Home evolved Node B (HeNB). Technical Specification TS 33.302 v11.2.0, 3G Partnership Project, June 2011.

[2] 3GPP. Universal Mobile Telecommunications System (UMTS); LTE; Lawful interception requirements. Technical report, 3rd Generation Partnership Project, 2011. 3GPP TS 33.106 version 10.0.0 Release 10.

[3] 3GPP. UTRAN architecture for 3G Home Node B (HNB); Stage 2. Technical Specification TS 25.467 v10.2.0, 3G Partnership Project, June 2011.

[4] M. Basseville and I. V. Nikiforov. *Detection of Abrupt Changes: Theory and Application*. Prentice-Hall, Inc., 1993.

[5] L. E. Baum and G. R. Sell. Growth functions for transformations on manifolds. *Pacific J. Math*, 27:211–227, 1968.

[6] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[7] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander. Lof: Identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, 2000.

[8] X. Ding, L. He, and L. Carin. Bayesian robust principal component analysis. *Image Processing, IEEE Transactions on*, 20(12):3419–3430, Dec 2011.

[9] W. Enck, P. Traynor, P. McDaniel, and T. La Porta. Exploiting Open Functionality in SMS-Capable Cellular Networks. In *Proceedings of the 12th ACM conference on Computer and communications security*, CCS '05, pages 393–404, New York, NY, USA, 2005. ACM.

[10] B. T. Y. X. X. N. F. Wang, H. Zhu and Y. Yang. A hmm-based method for anomaly detection. In *IEEE Conference on Broadband Network and Multimedia Technology*, pages 276–280, 2011.

[11] N. Golde, K. Redon, and R. Borgaonkar. Weaponizing Femtocells: The Effect of Rogue Devices on Mobile Telecommunications. In *Proceedings of the 19th Annual Network and Distributed System Security Symposium (NDSS)*, February 2012.

[12] GSMA. SMS Spam and Mobile Messaging Attacks Introduction, Trends and Examples. Online `http://www.gsma.com/technicalprojects/wp-content/uploads/2012/04/srssmsspamandmobilemessagingattacksthreatsandtrendswp.pdf`, Jauary 2011.

[13] GSMA. Handset theft. 2012.

[14] I. T. Jolliffe. *Principal Component Analysis*. Springer Verlag, 1986.

[15] S. S. Joshi and V. V. Phoha. Investigating hidden markov models capabilities in anomaly detection. In *ACM Proceedings of the 43rd Annual Southeast Regional Conference*, pages 98–103, 2005.

[16] S. Kent and K. Seo. Security Architecture for the Internet Protocol. RFC 3401, December 2005.

[17] M. Lange, S. Liebergeld, A. Lackorzynski, A. Warg, and M. Peter. L4Android: A Generic Operating System Framework for Secure Smartphones. In *Proceedings of the 1st ACM Workshop on Security and Privacy in Smartphones and Mobile Devices*, SPSM '11, 2011.

[18] P. P. C. Lee, T. Bu, and T. Woo. On the Detection of Signaling DoS Attacks on 3G Wireless Networks. In *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pages 1289–1297, Anchorage, AK, May 2007.

[19] S. Li and R. Schmitz. A novel anti-phishing framework based on honeypots. In *Proceedings of 4th Annual APWG eCrime Researchers Summit (eCRS'2009)*. IEEE, 2009.

[20] S. Liebergeld, M. Lange, and C. Mulliner. Nomadic Honeypots: A Novel Concept for Smartphone Honeypots. In *Proceedings of the Workshop on Mobile Security Technologies 2013*, Most 2013, 2013.

[21] I. L. MacDonald and W. Zucchini. *Hidden Markov and other models for discrete-valued time series*. Monographs on statistics and applied probability. Chapman & Hall, 1997.

[22] S. Munaut. IMSI Detach DoS. `http://security.osmocom.org/trac/ticket/2`, May 2010.

[23] E. S. Page. Continuous inspection schemes. *Biometrika*, 41:100–115, 1954.

[24] B. Pardo and W. Birmingham. Modeling form for on-line following of musical performances. In *National Conference on Artificial Intelligence*, volume 20, page 1018, 2005.

[25] L. Piyathilaka and S. Kodagoda. Gaussian mixture based hmm for human daily activity recognition using 3d skeleton features. In *IEEE Conference on Industrial Electronics and Applications*, pages 567–572, 2013.

[26] L. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77:257–286, Feb 1989.

[27] J. Serror, H. Zang, and J. C. Bolot. Impact of paging channel overloads or attacks on a cellular network. In *Proceedings of the 5th ACM workshop on Wireless security*, WiSe '06, pages 75–84, New York, NY, USA, 2006. ACM.

[28] G. Strang. *Linear Algebra and its Applications*. Harcourt Brace Jovanovich, 3rd edition, 1988.

[29] N. Times. ellphone Thefts Grow, but the Industry Looks the Other Way. Online `http://www.nytimes.com/2013/05/02/technology/cellphone-thefts-grow-but-the-industry-looks-the-other-way.html?pagewanted=all&_r=0`, 2013.

[30] A. R. Tom Ritter, Doug DePerry. I Can Hear You Now: Traffic Interception and Remote Mobile Phone Cloning with a Compromised CDMA Femtocell. `https://www.blackhat.com/us-13/briefings.html#Ritter`, August 2013.

[31] P. Traynor, M. Lin, M. Ongtang, V. Rao, T. Jaeger, P. Mcdaniel, and T. La Porta. On Cellular Botnets: Measuring the Impact of Malicious Devices on a Cellular Network Core. In *Computer and Communications Security*, pages 223–234, 2009.

[32] P. Traynor, M. Lin, M. Ongtang, V. Rao, T. Jaeger, P. McDaniel, and T. La Porta. On cellular botnets: measuring the impact of malicious devices on a cellular network core. In *Proceedings of the 16th ACM conference on Computer and communications security*, pages 223–234. ACM, 2009.

[33] W. Wong and M. Stamp. Hunting for metamorphic engines. *Journal in Computer Virology*, 2:211–229, 2006.

[34] W. H. Woodall. The use of control charts in health-care and public-health surveillance. *Journal of Quality Technology*, 38:89–134, 2006.

[35] M. Xie, T. N. Goh, and X. S. Lu. A comparative study of ccc and cusum charts. *Quality and Reliability Engineering International*, 14:339–345, 1998.

[36] S. Yang, K. Kalpakis, and A. Biem. Detecting road traffic events by coupling multiple timeseries with a nonparametric bayesian method. *Intelligent Transportation Systems, IEEE Transactions on*, 15(5):1936–1946, Oct 2014.